Self-fulfilling Prophecies in the Transition to Clean Technology

Sjak Smulders and Sophie Zhou*
July 4, 2025

Abstract

Technological lock-in has been a standard explanation for the slow take-off of clean innovation, but is hard to reconcile with forward-looking investors who anticipate the eventual switch to clean technologies. We provide an alternative explanation: strategic investment complementarities shape innovation and self-fulfilling prophecies can lead to delayed low-carbon transition. We analyze a standard directed technical change model with clean and dirty inputs. We find that when the two are good substitutes, two stable steady states can co-exist, each allowing multiple transitional paths. Optimal low-carbon transition requires a Pigouvian tax rule combined with a coordination device; commitment to a Pigouvian tax trajectory cannot solve a coordination failure.

JEL codes: C62, O33, O38, O44, Q55

Keywords: Directed technical change, strategic complementarity, multiple equilibria, delayed low-carbon transition, stranded asset

*Smulders: Tilburg Sustainability Center, Department of Economics and CentER, Tilburg University, P.O. Box 90153, Tilburg 5000 LE, The Netherlands. j.a.smulders@tilburguniversity.edu. Zhou: Department of Economics, University of Reading, PO Box 217, Reading, RG6 6AH, United Kingdom. s.l.zhou@reading.ac.uk. This paper was previously circulated under the title "Believe it or not: how directed technical change and a bit of optimism lead the economy towards green growth". We thank Peter K. Kruse-Andersen, Inge van den Bijgaart, Stefano Carattini, Guido Cozzi, Reyer Gerlagh, Bard Harstad, Rob Hart, Gilbert Kollenbach, Derek Lemoine, Linus Mattauch, Conny Olovsson, Pietro Peretto, Francesco Ricci, Florian Schütt, and Alex Teytelbaum for helpful discussions and comments. We would also like to thank participants at GSS seminar at Tilburg University, SURED Conference in Banyuls-sur-Mer, EAERE Conference in Athens, EEA-ESEM Conference in Lisbon, research seminars at Oxford University (2018), Oslo University (2019), 2020 Conference on New Developments in Business Cycle Analysis, research seminar at Arizona University (2021), 2022 Bundesbank Friendlyfaces Climate Workshop, webinar at ESCB Research Cluster on Climate Change (2022), and research seminar at Edinburgh University (2024) for comments and questions. This research was made possible by financial support of PBL Netherlands Environmental Assessment Agency.

1 Introduction

As ever more frequent natural disasters drive home the message of an urgent need for climate actions, climate risk is becoming a mainstream consideration. Anticipating more stringent climate policies and recognizing the inevitability of a transition towards cleaner production, more and more companies put forward their net-zero pledges (BEIS, 2021; Thorbecke, 2021). Nonetheless, policy makers are concerned about the slow take-off of private market funding for the transition (Draghi, 2024; Nerlich et al., 2025). The warnings of asset stranding (van der Ploeg and Rezai, 2020) highlight the concern that the private sector may not be transitioning to clean technologies promptly and optimally. The recent evidence of a declining share of clean energy innovation (Popp et al., 2022) serves as a case in point.

One standard explanation for the slow transition is path dependency and technological lock-in, as emphasized by the directed technical change (DTC) literature (see for example Acemoglu et al., 2012). Companies are reluctant to give up their firm position in polluting industries and forego historically established markets. However, while ignoring history is costly, ignoring the future might be even more so, particularly for long-lived assets and forward-looking decisions such as innovation and investments. In addition, the narrative of path dependency is at odds with the rise and fall of clean innovation, as documented by Probst et al. (2021).

Instead of path dependency, a coordination failure in technology investment may explain both the lack of investment and the peaking of clean innovation. As the early literature on coordination failures has pointed out in a different context (Krugman, 1991; Matsuyama, 1991), individual investors prefer to direct their investment to markets that are believed to be big and profitable in future, but the size of future markets depends on aggregate investments of these investors themselves. Consider for example, in the context of the energy transition, the decision to invest in the electric vehicle (EV) market versus the traditional car market. Each type of cars requires their type-specific inputs: e.g. motors and batteries for EVs and combustion engines and fuel for traditional cars. If the two car types are good substitutes from the user perspective, innovation in batteries benefits predominantly the EV sector and directly increases the profitability of electric motor innovation, much more than combustion engines innovation would do. These complementarities in firms' investment decisions create multiple equilibria. If all believe a particular market is the most profitable one to invest in, investors flock in this market and the resulting bigger market justifies their collective choice. But another market would have been targeted with different beliefs. Along this mechanism, self-fulfilling prophecies may strengthen brown investment and counteract green policies.

This paper is the first that includes coordination failures in a general equilibrium model of technology and climate dynamics. It develops a DTC model in which two goods, produced using clean and polluting inputs respectively, compete for market shares and revenues, while monopolistic firms supplying these inputs invest in R&D in order to increase their firm value. How much a single forward-looking firm benefits from innovation is measured by the discounted value of the induced future profits, which depend both on total spending in the sector (clean versus polluting) and on the firm's share in sectoral spending. The firm captures a share proportional the quality of the firm's product relative to average quality in the sector. Over time, spending shifts to the sector that innovates most since polluting and clean goods are gross substitutes. Investment by firms within a sector thus generates a demand externality. Innovation by other firms in the same sector also lowers one's own revenue share, leading to a business stealing effect. For sufficiently good substitutability, the demand externality outweighs the business stealing, which makes innovations by firms within a sector strategic complements: innovation by one firm raises the return to innovation for other firms in the sector, which in turn raises their own incentive to innovate.

As a first contribution, this paper analyzes how coordination problems arise in a standard DTC setting. Expectations or beliefs about future environmental-friendliness of innovation can overturn the lock-in in polluting technology and create self-fulfilling prophecies. In particular, we find that two stable steady-state equilibria can co-exist – one in which dirty technologies dominate and one in which only clean technologies are used – and that the selection between the two involves a coordination problem. In addition to the selection of the steady state, the transition towards a steady state is itself subject to coordination. We find that multiple transitional paths exist to each of the steady states that are consistent with rational deterministic expectations, i.e. expectations that can be rationalized in an equilibrium without stochastic shocks. Furthermore, some of the equilibrium paths involve multiple changes in the direction of innovation. In the context of the green transition, this implies that non-monotonic transition patterns involving rises and falls of clean innovation are possible.

As a second contribution, this paper integrates macro-economic investment coordination problems with climate dynamics and derives optimal dynamic policies. Even if climate policy can eliminate the undesirable steady state, it may not rule out delays and excess asset-stranding during the transition towards the clean steady state, as self-fulfilling prophecies continue to exist for the selection among multiple transitional paths. Our results are consistent with the observation that firms continue to invest in polluting technologies despite the inevitability of more stringent climate policy and the transition towards clean technologies. Our results also show that a Pigouvian carbon tax is insufficient for achieving the optimal low-carbon transition. The climate damage and the coordination problem cause two separate externalities that require two instruments. Under a Pigouvian tax rule, innovators internalize their impact on future climate damages, but a second instrument is needed to coordinate innovators on the appropriate transitional path, e.g. an R&D tax for polluting technologies. With investment complementarities, commitment to the first-best pollution tax level still leaves room for coordination on the wrong (i.e. polluting) innovation path. Strategic complementarities still enhance relative returns to innovation if investments are concentrated in these activities. Any tax that affects one sector more than another only overcomes the coordination failure if the tax level is big enough, depending on the strength of the investment complementarities. The required pollution tax level in general does not coincide with the required level to internalize the pollution damages. In the race between coordination failure and pollution externality, commitment to a pollution tax alone cannot set all incentives right. Our results thus lend support to the simultaneous use of multiple policy instruments rather than relying on a Pigouvian carbon tax alone.

Related literature Our paper contributes to several strands of literature. Firstly, our paper is related to endogenous growth models featuring multiple equilibria. Early papers in this area feature increasing returns in production with Marshallian externalities: the production of individual firms depend on the aggregate stock of physical capital (Benhabib and Farmer, 1994; Boldrin and Rustichini, 1994) or human capital (Benhabib and Perli, 1994). Scarcity of resource capital is central in van der Meijden and Smulders (2017), who study complementarity between resource conservation decisions by resource owners and investments in resource saving by innovators. In Gali (1996) complementarity originates from the response of the aggregate demand elasticities to changes in the composition of demand. Cozzi (2005), Cozzi (2007), and Gil (2013) point out that multiplicity arises in certain creative destruction endogenous growth models. In their focus on balanced growth paths, these papers deal with neither technology transitions in general nor the specific challenge of a transition to a carbon-neutral technology-based economy, which is the explicit topic of our paper.[1] Moreover, different from these papers, we stress demand externalities in input markets as a novel source of multiplicity in growth models.

Secondly, this paper analyzes investment incentives in the tradition of the DTC literature (Acemoglu, 2002; Acemoglu et al., 2012; Hassler et al., 2021), in which different factors of production (e.g. skilled versus unskilled labor, or clean versus polluting energy inputs) compete for innovation. If goods produced using these competing factors are gross substitutes, long-run innovation is concentrated in the sector with the least-cost factor supply (Hémous and Olsen, 2021). This drives lock-in in fossil inputs (Acemoglu et al., 2012), the potential of subsidies to green technologies to change the direction of innovation (Hart, 2019), and the possibility of rising extraction costs to induce the transition from fossil to clean (Lemoine, 2024). The existing literature concludes first-best and second-best instruments are needed to avoid costly lock-in or trigger the transition. Our analysis focuses on an hitherto overlooked aspect: gross input substitution may generate self-fulfilling prophecies in the direction of technical change, which changes the perspective on pollution tax and innovation subsidy policies.[2]

[1] Bretschger and Schaefer (2017)'s endogenous growth model builds on Gali (1996) and studies multiple equilibria in the built-up of green technology capacity, but without modelling the transition from polluting to clean technology as a transition process.

[2] Sturm (2023) also studies how coordination failures can be overcome in a dynamic setting. The setting differs: in his model the coordination failure is the only externality and investors move

The analysis deviates from the traditional assumption of one-period patents in this literature by assuming an infinite patent length. Longer-lived assets give a stronger role for expectations about future profitability, even though, as Section 3.3 shows, the existence of strategic complementarity does not necessitate a long patent length. Our paper thus contributes by pointing out that multiple equilibria can be an inherent feature of DTC models.

Thirdly, by reconciling the observation of continued innovation in polluting technologies with the expectation of an eventual switch to clean technologies, our paper contributes to the understanding of asset stranding (van der Ploeg and Rezai, 2020). Asset stranding is often attributed to the disorderly transition in which society faces uncertainty about the cost and distributional effects of green investments, and investors need to cope with uncertain policy revisions in the course of transition (Kalkuhl et al., 2020). Our paper is complementary by explaining asset stranding as a result of a suboptimal coordination of market expectations on little investment in green technologies. Our results also point to the limitation of using only a price instrument: even when environmental policy imposes a carbon tax that reflects future marginal damages or when markets price in a carbon risk premium, the economy may still coordinate on a suboptimal transition path.

In the rest of the paper, Section 2 presents the model, and Section 3 solves the decentralized equilibrium. Section 4 shows that multiple steady states and transitional paths exist. Section 5 solves the planner's problem and analyzes the implication of multiple equilibria for climate policy. Finally, Section 6 concludes.

2 The model

We start from the Acemoglu et al. (2012) model in continuous time and tailor it to the analysis of self-fulfilling prophecies through two modifications. First, patents last forever instead of one period only, which allows expected future profits and policies to affect innovation more generally. Second, labor is mobile between production and research, which allows expectations to not only affect the direction but

sequentially, i.e. his focus is on a sequential non-repeated interaction between investors while in this paper investors invest simultaneously and repeatedly.

also the intensity of innovation.[3]

2.1 Households

The representative household derives utility from consumption of the aggregate good, $C(t)$, which is a CES aggregate of the clean ($C_c(t)$) and dirty ($C_d(t)$) sectoral good:

$$C(t) = \left[C_c(t)^{\frac{\sigma-1}{\sigma}} + C_d(t)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{1}$$

where σ is the elasticity of substitution between the two sectoral goods. Consumers care about not only aggregate consumption but also climate damage according to the following instantaneous utility function

$$U(t) = \ln\left[(1 - D(S(t)))C(t)\right], \tag{2}$$

where the damage factor $D(S(t))$ is an increasing function of the carbon stock $S(t)$.

Households supply inelastically one unit of labor at wage $w(t)$, invest in assets $V(t)$ with return $r(t)$, and discount future utility at rate ρ. They maximize lifetime utility $W_0 = \int_0^\infty \ln[(1-D(S(t)))C(t)]e^{-\rho t}dt$ subject to the intertemporal budget constraint $\dot{V}(t) = r(t)V(t) + w(t) - P_c(t)C_c(t) - P_d(t)C_d(t)$, while taking climate damage $D(S(t))$ as given.

2.2 Final goods producers

There are two final goods sectors in the economy, clean ($Y_c(t)$) and dirty ($Y_d(t)$). Since the sectoral goods are substitutes, with constant elasticity $\sigma > 1$, total pollution intensity of the economy can be changed by substituting clean Y_c goods for polluting Y_d goods.[4] The sectoral goods are produced by competitive producers using

[3] These two modifications make the analysis richer, but neither labor mobility nor infinite patent length is needed to generate the qualitative results of our analysis. Online Appendix OA4 shows that the dynamic model with segmented labor market (as in Acemoglu et al. (2012)) is not analytically more tractable than the less restrictive model presented in the main text. Online Appendix OA3.2 relaxes the assumption of infinite patent length. Further, in Section 3.3 we show that complementarities and multiple equilibria arise in a static one-period setting.

[4] Labor and intermediates enter into the production of clean and polluting goods, which then enter production of the aggregate consumption good. This is isomorphic to a setting with a homogenous consumption good, C, being produced with clean inputs C_c and polluting inputs C_d. Hence σ takes the role of input substitution, i.e. substitution between clean and dirty inputs. This is also

labor and a continuum of sector-specific intermediates following a Cobb-Douglas technology:

$$Y_j(t) = L_j(t)^{1-\alpha} \int_0^1 q_{ji}(t) x_{ji}(t)^{\alpha} di, \qquad j \in \{c, d\} \tag{3}$$

where $L_{jt}(t)$ is the production labor hired in sector j, $q_{ji}(t)$ and $x_{ji}(t)$ are the quality and quantity of the intermediate good i in sector j. Final goods producers are price takers in both the final goods market and the factor markets. They decide on the factor demand ($L_{jt}(t)$, $\{x_{ji}(t)\}_{i=0}^{1}$) to maximize their profit $\pi_j(t) = P_j(t)Y_j(t) - w(t)L_j(t) - \int_0^1 P_{ji}(t)x_{ji}(t)di$, where $P_j(t)$ is the price of the final good j, $w(t)$ is wage, and $P_{ji}(t)$ is the price of the intermediate good i in sector j.

While the clean sector is carbon free, the production of the dirty sector generates a_d units of carbon emissions $E(t)$ per unit of output $Y_d(t)$:[5]

$$E(t) = a_d Y_d(t). \tag{4}$$

2.3 Intermediate goods producers

Each intermediate good $x_{ji}(t)$ is produced by a monopolist using the final goods of that sector. The unit cost of production increases with its quality $q_{ji}(t)$ so that one unit of intermediate good requires $q_{ji}(t)$ units of final goods j.[6] Monopolists hire

isomorphic to a more elaborate production structure with the aggregate good being a Cobb-Douglas aggregate of competitive sectors, each of which uses clean and dirty inputs with constant elasticity of substitution σ. This shows that parameter σ reflects substitution and pollution reduction options throughout the economy.

[5]This follows Acemoglu et al. (2012), but can be generalized in three respects without implications for our context. First, when Y_c production also pollutes, but relatively less than Y_d production, aggregate emissions reductions still work through the same channel, namely a shift away from the d sector triggered by environmental policy or directed innovation. Second, instead of fixing emissions per unit of Y_d, there could be abatement within the d sector. This would add another layer of abatement options which would modify only the strength of the economy-wide abatement channel currently in the model (i.e. through shifts between the sectors). Within-sector abatement becomes relevant mainly for policy design – we discuss this in section 5.3. Third, as in Acemoglu et al. (2012), more production of the Y_d good generates more emissions, even if the production expansion is driven by innovation. Thus innovation in the d sector implies pollution-using technical change, as the counterpart of pollution-saving technical change in the c sector, which captures documented innovation patterns (e.g. Dugoua and Gerarden, 2025). Innovation in the d sector captures technical change in the fossil extraction sector (Acemoglu et al., 2023) or quality improvement or cost reduction in any relatively polluting sector (Fried, 2018).

[6]While it is not only realistic that higher quality products are more costly to produce, this assumption also prevents that innovating firms become disproportionately large; technically, it ensures

research labor $s_{ji}(t)$ for R&D to improve the quality of their products according to:[7]

$$\dot{q}_{ji}(t) = \mu Q_j(t) s_{ji}(t), \tag{5}$$

where μ is a research productivity shifter and $Q_j(t) \equiv \int_0^1 q_{ji}(t)di$ is the sector-wide average quality level. The latter proxies for the knowledge stock[8] on which innovation builds: the more knowledge accumulated for a class of technology, the more there is to draw on in R&D.[9]

The intermediate goods monopolists choose the amount of production $x_{ji}(t)$ and the level of research effort $s_{ji}(t)$ to maximize the net present value of its profits $\pi_{ji}(t) = x_{ji}(t)(P_{ji}(t) - P_j(t)q_{ji}(t))$, subject to the demand for their goods as derived from the factor demand by the final goods producers.

2.4 Climate dynamics and damage

Following Golosov et al. (2014), we assume $1 - D(S(t))$ to be

$$1 - D(S(t)) = \exp\left(-\gamma(S(t) - \bar{S})\right), \tag{6}$$

where $\gamma > 0$ is a damage parameter and $\bar{S}$ is the pre-industrial carbon concentration.

Carbon concentration, in turn, increases with the emission from the dirty output production and is given by

$$S(t) = \int_{-\infty}^{t} \left[\phi_L + \phi_D e^{-\delta(t-s)}\right] E(s) ds, \tag{7}$$

where ϕ_L is the share of the carbon emission that stays in the atmosphere forever, $\phi_D/(1-\phi_L)$ is the share from the remaining emission that does not exit the atmosphere immediately but rather decays gradually, and δ is the geometric decay rate.

---

that only aggregate quality and not the distribution of quality in the sector determines sectoral spending.

[7] We choose to model innovation as the result of inhouse R&D à la Smulders and Nooij (2003), rather than "creative destruction". The former gives simpler mathematical expressions and seems to be at least equally empirically relevant as is shown recently by Garcia-Macia et al. (2019). This modeling choice does not alter the qualitative results in this paper, see Online Appendix OA5.

[8] We use "knowledge" and "technology" interchangeably throughout.

[9] In Appendix OA3 we show that positive spillovers in innovation are not crucial for our main result of self-fulfilling prophecies; even negative spillovers (e.g. when $Q_j(t)$ in (5) gets a negative exponent) are compatible with self-fulfilling prophecies, as long as they are not too strong and are offset by demand complementarities.

3 Direction of technical change in decentralized equilibrium

We derive the equilibrium conditions and show that investment complementarities drive self-fulfilling prophecies in the innovation decision. To highlight that this finding is independent of the savings decision, we defer the dynamic equilibrium analysis of the labor market to Section 4. Time arguments are omitted whenever no confusion would arise.

3.1 Equilibrium conditions

Producers Profit maximization of the final goods producers gives the usual factor demand, for labor and intermediates, respectively:

$$w = (1-\alpha)P_j \frac{Y_j}{L_j}, \tag{8}$$

$$P_{ji} = \alpha P_j L_j^{1-\alpha} q_{ji} x_{ji}^{\alpha-1}. \tag{9}$$

Intermediate goods monopolists maximize the net present value of their profits subject to the demand for their goods given by (9), leading to the following first order conditions (see Appendix OA1 for detailed derivation):

$$P_{ji} = \frac{1}{\alpha} P_j q_{ji}, \tag{10}$$

$$\mu Q_j \lambda_{ji} \leq w \perp s_{ji} \geq 0, \tag{11}$$

$$\dot{\lambda}_{ji} = r\lambda_{ji} - \frac{\partial \pi_{ji}}{\partial q_{ji}}, \tag{12}$$

where λ_{ji} denotes the firm's shadow value of quality improvements. Equation (10) is the usual markup rule. In (11), $\mu Q_j \lambda_{ji}$ represents the contribution of a marginal unit of research to the present value of the firm's future profit, while wage w captures the marginal cost of research effort. This equation characterizes the monopolist's investment decision by equality of marginal benefits and costs in case of active research. Finally, (12) is the arbitrage equation that determines the shadow value of quality improvements.

Combining demand (9) and markup rule (10), we conclude that all firms in the same sector sell the same equilibrium quantity. Next, calculating their profits we find that these are linear in their own good's quality and that all monopolists enjoy the same marginal profit from quality improvement $\partial\pi_{ji}/\partial q_{ji}$. Thus, from (12), they face the same shadow value of innovation λ_{ji}; we henceforward write λ_j. The forward-looking component of a firm's investment decision thus depends on sector-level variables only. We can characterize the production side of the economy as:

$$Y_j = \alpha^{\frac{2\alpha}{1-\alpha}} Q_j L_j, \tag{13}$$

$$\dot{Q}_j = \mu s_j Q_j, \tag{14}$$

$$\pi_{ji} = (1-\alpha)\alpha P_j Y_j \frac{q_{ji}}{Q_j}, \tag{15}$$

where $s_j \equiv \int_0^1 s_{ji} di$ is the sectoral research input. According to (13) and (14), sectoral production reduces to a Ricardian production function in which productivity is proportional to the aggregate technology stock Q_j, which grows with aggregate research effort in that sector. All firms in the sector sell the same quantity, but firms with higher quality can sell at a higher price, (10), and thus make higher profits. Since final goods firms spend a constant fraction on intermediate goods and intermediate goods producers have a constant markup, total profits are a constant share $(1-\alpha)\alpha$ of revenue, while an individual firm's share in total profits is q_{ji}/Q_j according to (15).

Households Utility maximization of the households leads to the usual static demand functions and the Euler equation:

$$\frac{C_c}{C_d} = \left(\frac{P_c}{P_d}\right)^{-\sigma}, \tag{16}$$

$$r - \hat{P} = \rho + \hat{C}, \tag{17}$$

where P is the price index of consumption defined by $PC = P_c C_c + P_d C_d$ and we use hats to denote growth rates, $\hat{x} \equiv \dot{x}/x$ for any x. Equation (16) shows that relative demand responds to relative price with elasticity σ. (17) shows that households require a real rate of return $(r - \hat{P})$ on their savings that reflects their impatience (ρ) and a premium for postponing consumption $(\hat{C})$.

Market clearing Goods market clearing requires that in each sector total production net of intermediate input use equals consumption. Using equilibrium quantities of x_{ji} and Y_j (see (13)), we find

$$C_j = Y_j - \int_0^1 q_{ji} x_{ji} di = (1-\alpha^2) Y_j. \tag{18}$$

Labor market clearing requires that total (exogenous) supply equals demand for production and research:

$$1 = L + s_c + s_d, \tag{19}$$

where $L \equiv L_c + L_d$ is the total production labor.

Static equilibrium The allocation of profits, production labor, and consumer spending over the clean and dirty sectors depend on the pre-determined (i.e. history dependent) technology stocks Q_c and Q_d. We define θ_j to reflect the relative importance of a given technology stock Q_j as:

$$\theta_j \equiv \frac{Q_j^{\sigma-1}}{Q_c^{\sigma-1} + Q_d^{\sigma-1}}. \tag{20}$$

From equations (8), (13), (16), and (18), we find

$$\frac{L_c}{L_d} = \frac{P_c C_c}{P_d C_d} = \left(\frac{Q_c}{Q_d}\right)^{\sigma-1} = \frac{\theta_c}{\theta_d}, \tag{21}$$

Since $\theta_c + \theta_d = 1$, in equilibrium θ_j equals sector j's labor and spending shares:

$$\theta_j = \frac{L_j}{L} = \frac{P_j C_j}{PC}. \tag{22}$$

3.2 Investment complementarities

From (15), profits of intermediate goods firms are related to spending shares:

$$\pi_{ji} = \frac{q_{ji}}{Q_j}(1-\alpha)\alpha\theta_j PY, \tag{23}$$

which shows that (i) total profits in sector j are proportional to spending $\theta_j PY$ and (ii) firms earn a share in total profits that reflects their relative quality q_{ji}/Q_j. Since profits are linear in own quality,[10] marginal profits depend on sectoral variables only: they are proportional to θ_j/Q_j. Using (20) and (21), we express relative marginal profits as:

$$\frac{\partial \pi_{ci}/\partial q_{ci}}{\partial \pi_{di}/\partial q_{di}} = \left(\frac{Q_c}{Q_d}\right)^{\sigma-2}. \tag{24}$$

Investment shifts spending shares and profits. A firm investing in higher quality directly experiences higher profits (through own quality q_{ji}), but non-investing firms are affected as well (through Q_j). As (21) shows, the relative revenue of the two sectors $(P_cC_c/(P_dC_d))$ increases with relative quality (Q_c/Q_d) if the two sectors are gross substitutes $(\sigma > 1)$. Higher quality in sector j shifts expenditure to this sector as long as $\sigma > 1$ and the sector as a whole reaps more profits (demand shift effect). Within the same sector, (23) shows that the relative profit of two firms $(\pi_{ji)}/\pi_{ji'})$ is equal to the relative quality of their own goods $(q_{ji}/q_{ji'})$. Non-investing firms see their share in these profits fall as their quality is relatively lower (business stealing effect). From (24) it is clear that the demand shift effect dominates the business stealing effect if and only if $\sigma > 2$, so that innovations by firms within the same sector become strategic complements: investment by any firm increases the profits from investment for all other firms in the sector.

In Appendix OA3, we generalize the model to allow for more sources of complementarities.[11] In particular, we show that complementarities become stronger with learning-by-doing spillovers in production (parameterized by ε) and with the share of own sector inputs in the production of intermediates (parameterized by $\omega \in [0,1]$). Cross-sectoral knowledge spillovers (parameterized by χ), on the other hand, do not affect complementarities at all. The condition for strategic complementarity is then $\sigma > 1 + (1-\omega\alpha)/((1-\alpha)(1+\varepsilon))$, with the model in the main text of this paper being a special case with $\varepsilon = 0$ and $\omega = 1$.

[10]Because the size of each firm is negligible relative to the sector, Q_j is considered exogenous by firms.

[11]Appendix OA3 replaces L_j in (3) by $Q_j^{\varepsilon}L_j$ to allow for production spillovers, replaces Q_j in (5) by $Q_j^{\eta+\chi}Q_{-j}^{\chi}(Q_c+Q_d)^{1-\eta-2\chi}$ to allow for intersecoral knowledge spillovers, and P_j in (10) by $P_j^{\omega}P^{1-\omega}$ to allow for intersectoral input-output linkages.

3.3 Coordination in the static setting: a thought experiment

As is well known, strategic complementarities may lead to coordination problems. This is immediate in a static setting. As a thought experiment, suppose our model economy exists only for one period and can employ s innovators. Innovation decision then becomes a static two-stage coordination game. From (5), innovation effort s_{ji} determines individual and sectoral technology stocks according to $q_{ji} = q_{ji0} + \mu s_{ji} Q_{j0}$ and $Q_j = (1+\mu s_j) Q_{j0}$. The marginal benefit of innovation for an individual firm is then given by $\partial \pi_{ji} / \partial s_{ji} = (\partial \pi_{ji} / \partial q_{ji}) \mu Q_{j0}$, while the marginal cost of innovation, the wage paid to scientists, is the same for all firms. From (24), the relative benefit of innovating in the clean sector is thus $(Q_c/Q_d)^{\sigma-2}(Q_{c0}/Q_{d0}$, while the relative cost is 1. If all firms together invest in clean only, i.e. if $s_c = s, s_d = 0, Q_c = Q_{c0}(1+\mu s), Q_d = Q_{d0}$, the relative benefit of clean innovation is $(1+\mu s)^{\sigma-2}(Q_{c0}/Q_{d0})^{\sigma-1}$; this investment is an equilibrium if this expression exceeds the unitary relative cost (and is a stable equilibrium if $\sigma > 2$). Similarly, all firms investing in dirty is a stable equilibrium if $(1+\mu s)^{\sigma-2}(Q_{d0}/Q_{c0})^{\sigma-1} > 1$. Hence if $(1+\mu s)^{-(\sigma-2)} < (Q_{c0}/Q_{d0})^{\sigma-1} < (1+\mu s)^{\sigma-2}$, both equilibria exist.[12] In other words, if the economy starts with the initial quality levels not too far apart, two stable equilibria exist. Following Krugman we call this range of initial conditions for which multiple equilibria exist the “overlap”.

The same logic holds in Acemoglu et al. (2012), since the innovation decision in their model has the same two-stage feature with a similar cost-and-benefit structure due to one-period patents and a similar production structure. Innovators only care about immediate profits and the economy evolves recursively. Acemoglu et al. (2012) acknowledge the multiplicity of equilibria in their appendix, but restrict their analysis to initial conditions and policies that make the equilibrium unique. The next section shows the coordination issue with full dynamics.

3.4 Multiple rational expectations equilibria

Back to our infinite horizon model, firms' innovation decisions are driven by the contribution of higher quality to *all* future profits, λ_j. Together with the research

[12] A third (unstable interior) equilibrium exists with some firms investing in the clean sector and other investing in the dirty sector.

productivity, μQ_j, it determines the value of innovation for future profits in sector j. We define the following measure of future market conditions in sector j:

$$m_j \equiv \frac{\mu Q_j \lambda_j}{\mu Q_c \lambda_c + \mu Q_d \lambda_d} = \frac{Q_j \lambda_j}{Q_c \lambda_c + Q_d \lambda_d}, \tag{25}$$

where $\mu Q_j \lambda_j$ is a sector-j firm's marginal value of innovation (see (11)) and m_j represents the (marginal) market valuation share of investing in that sector.[13] The forward-looking variable m_c reflects how promising it currently is to invest in future green market rather than the other market.

Equation (11) shows that innovation will not take place in the sector in which research labor generates the lowest expected value. Expressed in terms of the market valuation share, this condition reads:

$$m_j > 1/2 \;\; \Leftrightarrow \;\; \hat{Q}_j \geq \hat{Q}_{-j} = 0. \tag{26}$$

This inequality delineates three innovation regimes: if $m_c > 1/2$ the economy is in the *clean-only innovation regime* with $s_c \geq 0$ and $s_d = 0$; if $m_c < 1/2$ the economy is in the *dirty-only innovation regime* with $s_c = 0$ and $s_d \geq 0$; the *simultaneous-innovation regime* with $s_c \geq 0$ and $s_d \geq 0$ requires $m_c = 1/2$. Market valuation share m_c is a continuous variable that cannot jump (unless unexpected shocks arise). Hence, if it starts above (or below) $1/2$ and the economy is in the clean-only (or in the dirty-only) regime, it remains so for some time. For consistency, we thus also require the simultaneous regime to be one where research is active in both sectors for a non-degenerate period of time. That is, the simultaneous regime is one where not only $m_c = 1/2$ is required but also $\dot{m}_c = 0$.

Under rational expectations, the forward looking valuation share m_c is connected to future development of market share θ_c, as it drives profits. The arbitrage equation (12) establishes this connection and implies:[14]

$$m_j > \theta_j \;\; \Leftrightarrow \;\; \hat{\lambda}_j > \hat{\lambda}_{-j}. \tag{27}$$

Intuitively, if the market valuation share of a sector is larger than its current market

[13] The direction of innovation is guided by the relative marginal innovation value $\lambda_c Q_c / \lambda_d Q_d$, which the variable m_c maps to the [0,1] interval so that it is directly comparable to θ_c.

[14] From (12) we find $\hat{\lambda}_c > \hat{\lambda}_d \Leftrightarrow [(\partial \pi_{ci}/\partial q_{ci})/\lambda_c]/[(\partial \pi_{di}/\partial q_{di})/\lambda_d] < 1$. From (24), (21), and (25), this is equivalent to $\theta_c m_d / \theta_d m_c < 1$, or $m_c > \theta_c$.

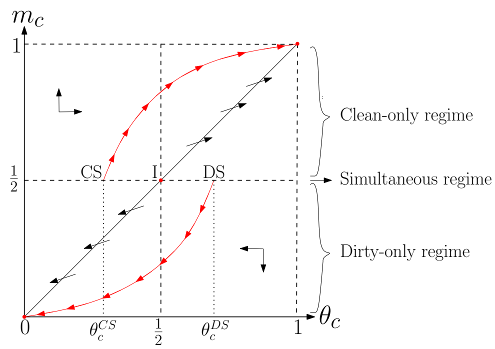

Figure 1: The (θ_c, m_c) equilibrium dynamics for $\sigma > 2$, showing the direction of technical change (small arrows) and the overlap $[\theta_c^{CS}, \theta_c^{DS}]$.

share, investors anticipate a growing market for that sector and the value of quality improvement grows faster in that sector than in the other sector.

To see how the direction of technical change evolves over time, we time differentiate (20) and (25):

$$\dot{m}_c = m_c(1-m_c)\left(\hat{\lambda}_c - \hat{\lambda}_d + \hat{Q}_c - \hat{Q}_d\right) \tag{28}$$

$$\dot{\theta}_c = \theta_c(1-\theta_c)(\sigma-1)(\hat{Q}_c - \hat{Q}_d). \tag{29}$$

Equations (26)-(28) summarize the investment block of the model with the exception that the total research effort (i.e. the speed of innovation), $1-L$, is yet unknown, which depends on labor market equilibrium and the savings decision. Nonetheless, (26)-(28) allow us to analyze the direction of technical change independently of the rest of the model. Intuitively, the *relative* reward of investment pins down the direction of technical change for any speed of innovation.

We define a candidate steady state as an equilibrium in which θ_c and m_c are constant or asymptotically constant. They are steady states of the model conditional on yet undetermined savings decision $(1-L)$.[15] From (26)-(28), the phase diagram in Figure 1 and Proposition 1 follow.[16] The proofs for this and all subsequent propositions and lemmas are provided in Appendix A.

[15]Since (28) and (29) are part of the full model, any steady state of the full model requires the subsystem to be in steady state.

[16]Figure 1 is a two-dimensional projection of the full phase diagram, which is essentially the full phase diagram conditional on the savings decision $(1-L)$.

Proposition 1. *Suppose $\sigma > 1$. There exist three candidate steady states: a clean candidate steady state with $m_c = \theta_c = 1$ and innovation only in the clean sector, a dirty candidate steady state with $m_c = \theta_c = 0$ and innovation only in the dirty sector, and an (unstable) interior candidate steady state with $m_c = \theta_c = 1/2$ and innovation in both sectors. If and only if $\sigma > 2$, there exists a range of initial clean market share $\theta_{c,0}$, for which both the clean and dirty candidate steady states can be reached in a rational expectations equilibrium.*

In Fig. 1, the horizontal line $m_c = 1/2$ is the regime border between clean and dirty innovation only. The red dots indicate the candidate steady states; the small arrows give the dynamics consistent with (26)-(28). It follows (as detailed in the proof) that the path to the clean steady state slopes upward and, if and only if $\sigma > 2$, it locates above the 45 degree line. As a result there is an equilibrium path towards the clean steady state that starts to the left of point I. Symmetrically, an equilibrium path to the dirty steady state starts to the right of point I and is below the 45 degree line if and only if $\sigma > 2$. The red path in the upper (lower) part of the figure illustrates the former (latter) path. Define θ_c^{CS} (θ_c^{DS}) as the value for θ_c at which the clean (dirty) path intersects with the regime border. The clean (dirty) candidate steady state can be reached whenever the historically inherited market share starts above θ_c^{CS} (below θ_c^{DS}). These threshold values define the overlap: if the initial market share is in the overlap $[\theta_c^{CS}, \theta_c^{DS}]$, both steady states can be reached.

To understand the overlapping innovation paths, recall that future market size justifies current investment. With high substitution between the two goods, consumers relatively easily substitute towards the good that becomes relatively cheap. If the clean sector invests even when it is currently small, prices fall and consumers shift spending to the sector, leading to a larger future market size for clean goods. Thus, innovators realize a higher return on innovation in this market, further lowering prices and reinforcing the incentive for consumers to spend more on clean goods. The spending share of the sector keeps increasing so that ultimately spending on clean goods dominates the market. This explains why the economy can move to the clean steady state even starting from relatively small productivity.

The investment complementarity is independent of the knowledge spillovers implied by (5). Appendix OA3 shows that the overlap remains when knowledge spillovers are muted or when intersectoral knowledge spillovers are introduced. The reason is that research productivity is only affected by inherited knowledge stock, that is, knowledge stock accumulated up to that point. Since this stock is predetermined, knowledge spillover from current investment by rival firms cannot affect one's value of investment and is not a source of investment complementarity in the market. In contrast, through the demand externality total investment affects future profits and thus the value of investment.

The implication of Proposition 1 is that investors' beliefs matter. If market shares are within the overlap, innovators' beliefs about future profitability determine which of the multiple rational equilibria is chosen. If the current green share is small but still within the overlap and firms are optimistic about future green market size, they invest and their beliefs come true - a self-fulfilling prophecy. But if all firms have pessimistic beliefs about a green transition, all firms invest in brown and their beliefs are also rational.

4 Self-fulfilling prophecies in the market economy

4.1 General equilibrium

We now analyze the full dynamics of the model by including labor market equilibrium and households' savings decisions which together determine the allocation of labor over production and research.

Equilibrium in the labor market requires that the wage in production equals the opportunity cost in research and that the labor market clears – equations (8), (11), and (19). Equilibrium in the capital market requires that the rate of return to innovation equals the required rate of return by households – equations (12), (15), and (17). Together with (26)-(28) this defines the general equilibrium dynamics:

Lemma 1. *The dynamic equilibrium is characterized by the evolution of market valuation share m_c, clean market share θ_c, and production labor L:*

$$\dot{m}_c = m_c(1-m_c)\mu(s_c - s_d) + \alpha\mu L m_k(m_c - \theta_c) \tag{30}$$

$$\dot{\theta}_c = \theta_c(1-\theta_c)(\sigma-1)\mu(s_c - s_d) \tag{31}$$

$$\dot{L} = L\left[\alpha\mu\theta_k L - \mu s_k - \rho\right], \tag{32}$$

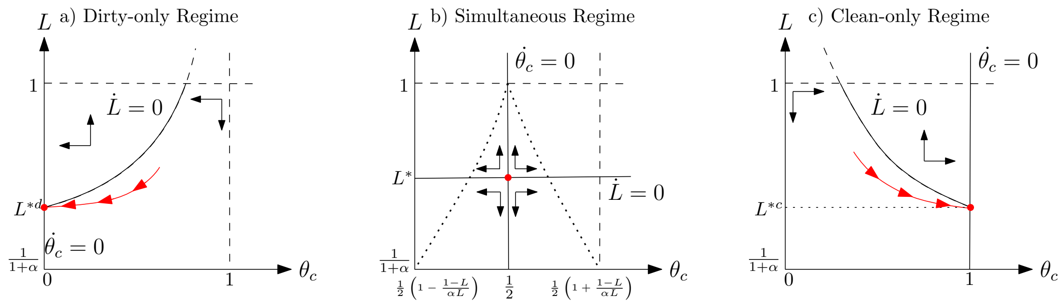

Figure 2: The (θ_c, L) phase diagrams (red dots mark steady states; the dotted lines in panel b mark the active-research bounds (33)). The dynamics of total innovation effort $s_c + s_d = 1 - L$ can be directly inferred for each regime.

where k denotes the research-active sector. In the clean-only regime ($k = c$): $m_c > 1/2$, $s_c = 1 - L$, and $s_d = 0$; in the dirty-only regime ($k = d$): $m_c < 1/2$, $m_k = 1 - m_c$, $\theta_k = 1 - \theta_c$, $s_c = 0$, and $s_d = 1 - L$; and in the simultaneous regime ($k = c, d$): $m_c = 1/2$, $s_c = (1 - L)/2 + \alpha L(\theta_c - 1/2)$, and $s_d = 1 - L - s_c$. Active innovation in both sectors in the simultaneous regime further requires

$$\frac{1}{2}\left(1 - \frac{1-L}{\alpha L}\right) \leq \theta_c \leq \frac{1}{2}\left(1 + \frac{1-L}{\alpha L}\right). \tag{33}$$

The lemma shows how the dynamics of market shares m_c and θ_c are governed by the dynamics of the labor allocation over production and innovation, (32). This general-equilibrium version of the Ramsey rule shows that the more the discount rate falls short of the real rate of return on innovation, the more is production of consumption goods postponed to the future (i.e. $\dot{L}/L > 0$). The return on innovation is increasing in market size $\theta_k L$, markup rate α, and R&D productivity shifter μ.

Lemma 1 makes clear that within each innovation regime, $\dot{\theta}_c$ and $\dot{L}$ do not depend on m_c. Accordingly, the dynamics of spending share θ_c and production labor L are represented in two-dimensional phase diagrams, one for each regime (see Figure 2). The vertical axis also determines labor for innovation, $1 - L = s_c + s_d$, which governs the *rate* of innovation. The *direction* of innovation continues to be determined by changes in the market valuation share m_c, which according to (30) depends on all three variables.

We define a steady state (an asymptotic steady state) as an equilibrium in which θ_c, m_c, and L are constant (asymptotically constant). We indicate (asymptotic) steady state values with an asterisk. The phase diagrams in Figures 1 and 2 identify the steady states as the intersection of the zero-motion loci.

Proposition 2. *Suppose $\sigma > 1$ and $\rho < \alpha\mu/2$. The decentralized economy has three steady states: (1) an unstable interior steady state with simultaneous R&D in both sectors, where $m_c^* = \theta_c^* = 1/2$, and $L^* = (1+\alpha)^{-1}(1+2\rho/\mu)$; (2) a saddlepath stable, asymptotic steady state with innovation in the dirty sector only, where $m_c \to m_c^{*d} = 0$, $\theta_c \to \theta_c^{*d} = 0$, $L \to L^{*d} = (1+\alpha)^{-1}(1+\rho/\mu)$; and (3) a saddlepath stable, asymptotic steady state with innovation in the clean sector only, where $m_c \to m_c^{*c} = 1$, $\theta_c \to \theta_c^{*c} = 1$, $L \to L^{*c} = (1+\alpha)^{-1}(1+\rho/\mu)$.*

In all three steady states, production occurs in both sectors. Although in the corner steady states one sector dominates with a market share approaching one, the absolute size of the dominated sector is non-zero and needs not be small. Consequently, all three steady states feature non-zero emissions.

The three steady states are illustrated by the red dots in Figure 2. The clean (dirty) steady state can only be reached from the clean-innovation (dirty-innovation) regime. Within the clean- or dirty-innovation regime, the red lines with arrows illustrate the unique path towards each of the two corner steady states. These paths are the projection of the three-dimensional saddlepaths on the (θ_c, L) plane, just as the red lines in Figure 1 depict the projection on the (θ_c, m_c) plane.

Proposition 2 focuses on the case where innovation is active (that is, $L < 1$) in all steady states. This requires $\rho < \alpha\mu$ for the corner steady states and $\rho < \alpha\mu/2$ for the interior steady state.

4.2 Self-fulfilling prophecies

To analyze the transition to the steady states, we first define equilibrium paths.

Definition 1 (Equilibrium Path)**.** *An equilibrium path is a sequence of $(\theta_c(t)$, $m_c(t)$, $L(t)$, $s_c(t)$, $s_d(t))$ that satisfies (30), (31), and (32) at any point in time, and approaches one of the two saddlepath-stable steady states as $t \to \infty$.*

An equilibrium path defined above is a sequence of $(\theta_c(t)$, $m_c(t)$, $L(t)$, $s_c(t)$, $s_d(t))$ that jointly maximizes firm profits and household lifetime utility, while clearing the factor, goods, and capital markets at the same time. Put in the context of rational expectations, an equilibrium path is the outcome of coordinated beliefs among all agents about all future actions of other agents and the resulting state of the economy. As agents are symmetric and atomistic, in equilibrium all agents must share the same belief, which can be represented by the equilibrium sequence of $(\theta_c(t),\ m_c(t),\ L(t),\ s_c(t),\ s_d(t))$. When agents coordinate on such a belief, it will turn out to be consistent with the optimizing behavior of all agents and market clearing of all markets at any point in time and thus consistent with their belief. Multiplicity arises if for the same initial condition multiple such equilibrium paths exist. In this case, the economy is free to select any such path, meaning that agents can coordinate on any such beliefs that will turn out to be rational. In our deterministic setting of the model, once a selection is made at time 0, there is no more uncertainty and rational expectations collapse into perfect foresight.

We further refer to any equilibrium path that approaches the clean steady state a clean path, and one that approaches the dirty steady state a dirty path.

Clean vs dirty path Proposition 2 confirms that the three candidate steady states listed in Proposition 1 are indeed steady states of the full model. As a result, the complementarity condition in Proposition 1 applies. Under complementarity ($\sigma > 2$), a range of initial conditions ("overlap") exists for which both the clean and dirty steady states can be reached.

We conclude that coordination of beliefs determines to which of the two steady states the economy evolves if it starts within the overlap. The full dynamic representation in Lemma 1 solves for the transition path and size of the overlap, as numerically illustrated in Figure 3 and analytically characterized in the following proposition:

Proposition 3. *Suppose $\sigma > 1$ and $\rho < \alpha\mu/2$. If and only if $\sigma > 2$, a range of initial states exists from which both corner steady states can be reached through a rational expectations equilibrium path. This range ("overlap") increases with substitutability σ and decreases with impatience ρ.*

The bigger the elasticity of substitution, the more consumers shift expenditure to the innovating sector and increase profits in this sector. Expectations of a bigger future market with associated investment in this market thus become more quickly justified. This explains why the overlap increases with the elasticity of substitution σ. An increase in impatience has the opposite effect: it reduces the present value of future profits and lowers investment incentives, thus decreasing the speed at which either sector can gain market size. If investors are infinitely impatient, only current profit matters and the investment becomes a static coordination game as in Section 3.3. In Appendix OA3.2, we further show that an increase in patent length has a similar effect to lowering impatience, as both increase the time horizon of investors.

Fast vs delayed transition Figure 1 shows that when tracing the saddlepathh from a corner steady state backward in time, the path eventually intersects with the regime border. At that point a regime switch must occur, which is also associated with the switch in direction of change for the spending share θ_c. This implies that when a saddlepath crosses the regime border, there are spending share levels that the saddlepath towards the same steady state visits multiple times. Thus, for some initial spending share, there are multiple equilibrium paths reaching the same steady state but each corresponding to a different length of transition time.

Since each such starting point is associated with a different future market condition $m_{c,0}$, the selection again depends on firms' beliefs of future profitability and thus the innovation decision of other firms in the economy. If the market expects ultimately a clean steady state, no matter in which innovation regime the economy starts, it must enter the clean-only regime before reaching the steady state. That is, the only rational expectation about a clean steady state is that, in the long run, all innovation will occur in the clean sector. For the short and medium run, however, if firms expect a speedy transition, it is only rational to invest only in the clean sector; if firms expect a delayed transition, it can be rational to temporarily invest in the dirty sector. Since firms' innovation decisions are strategic complements, such beliefs about the speed of transition can also be self-fulfilling.

Proposition 4. *From an initial $\theta_{c,0} \in \left[\theta_c^{CS}, \theta_c^{DS}\right]$, there are multiple transition paths towards each corner steady state. In particular,*

1. *a fast transition, in which the economy selects immediately the corner stable path, is always possible;*

2. *delays in transition before the economy finally selects the corner stable path are also possible; further, there exist* $\bar{\sigma}$ *and* $\bar{\bar{\sigma}}$*, where* $2 < \bar{\sigma} < \bar{\bar{\sigma}}$*, such that*
   (a) *if* $2 < \sigma < \bar{\sigma}$*, temporary simultaneous R&D is the only possible delay;*
   (b) *if* $\sigma \in [\bar{\sigma}, \bar{\bar{\sigma}})$*, delay must include temporary regime switches between the two corner innovation regimes;*
   (c) *if* $\sigma > \bar{\bar{\sigma}}$*, delay must include temporary stagnation with no R&D.*

The importance of σ for the pattern of delays again reflects the role of market responsiveness. For a relatively low σ, relative demand responds to innovation-induced changes in relative price only sluggishly. No matter which sector is expected to eventually dominate, changes towards this sector can occur only slowly. It takes a lot of cumulative changes to finally tip the balance towards one of the two sectors decisively. As σ becomes higher, the speed of demand responses picks up. An initial bias of expectations in favor of one of the two sectors can now be much more easily confirmed by market responses, and become self-reinforcing. This higher speed of demand response, however, also adds to volatility, as it becomes easier to tip the balance in favor of a specific sector. For very high value of σ, strong demand response means that a sector can be favored to dominate in the long run even if it initially lags far behind, as reflected by the large overlap size. Once the balance tips towards the laggard sector decisively, the expected demand in the initially leading sector is too low to incentivize any innovation in that sector. However, due to the small technology stock, productivity in innovation is also too low in the laggard sector and labor can be more productively used in production. Stagnation occurs as a result.

The implication of Proposition 4 is that even if firms believe that the economy will eventually transition to a clean steady state, how fast this happens is subject to coordination. If firms believe that there will be some period of indecisiveness concerning the relative profitability, it can indeed be rational to innovate only in the dirty sector temporarily. We thus offer a *forward-looking* explanation why we may observe predominantly dirty innovation, even if firms believe that the economy will switch to clean technologies in some future time, and why the green transition may exhibit non-monotonic patterns.

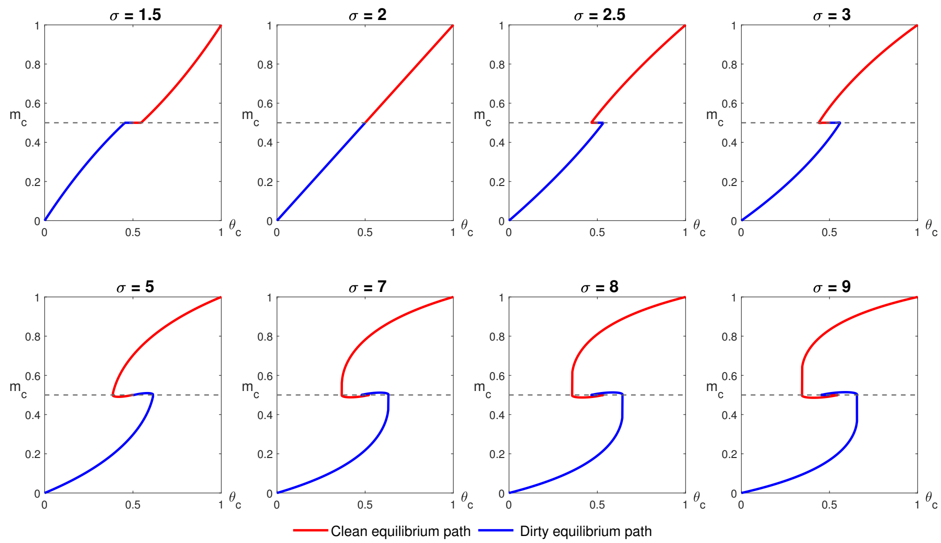

Figure 3: Equilibrium paths with different σ values

Numerical example To provide a graphical illustration of the overlap and the different transition scenarios, we now consider a numerical example. We set the labor share in production to the common value of two-third so that $\alpha = 1/3$ and the long-run growth rate to 1.25 percent. We set $\rho = 0.01$ and derive μ from the long-run growth rate. We allow the elasticity of substitution σ to vary.[17]

Figure 3 shows the projection of the equilibrium paths on the (θ_c, m_c) plane. As in Proposition 3, the clean and dirty equilibrium paths overlap for all σ values larger than 2. Further, the overlap region grows larger with increasing σ. In terms of the possible delays, for $\sigma = 2.5$ and $\sigma = 3$, the delays involve simultaneous R&D, as illustrated by the flat part of the equilibrium paths, where $m_c = 1/2$. With σ between 5 and 7, the equilibrium paths involve regime switches between the clean-only and dirty-only regimes. Finally, with $\sigma = 8$ and $\sigma = 9$, the equilibrium paths contain vertical sections corresponding to temporary stagnation, where $\dot{\theta}_c = 0$. These last equilibrium paths also contain multiple regime switches, featuring rises and falls of clean innovation along the transition.

[17] As can be seen from eq. (29), if $\sigma \in (0,1)$, the interior steady state is saddlepath stable, while the two corner steady states are unstable. By construction, no overlap can arise if $\sigma \in (0,1)$ and we thus ignore this case.

5 Self-fulfilling prophecies and optimal policy

We now characterize the social optimum and study the implications of coordination failure for optimal policy.

5.1 Social optimum

The planning problem is equivalent to maximizing (2) subject to (1), (3), (18), (19), (5), the equation of motion for carbon concentration $\dot{S}$ (derived from time differentiating (7)), and the non-negativity constraints of L_j and s_{ji} for all j and all i. The detailed solution to the planning problem is provided in Appendix OA2.

For the static optimum, the social planner corrects both the misallocation caused by the market power of the intermediate monopolists and that caused by the climate externality. Let τ_d^s denote the marginal cost of carbon emissions in terms of the dirty goods. The planner's sectoral output is given by

$$Y_c = \alpha^{\frac{\alpha}{1-\alpha}} Q_c L_c, \quad Y_d = ((1-\tau_d^s)\alpha)^{\frac{\alpha}{1-\alpha}} Q_d L_d. \tag{34}$$

The marginal cost of carbon emissions equals $\tau_d^s \equiv a_d \frac{\partial \dot{S}}{\partial E} \lambda_S / \zeta_d$, i.e. the social cost of the carbon stock in utility terms, λ_S, corrected for the effect of dirty production on the carbon stock $a_d \frac{\partial \dot{S}}{\partial E}$ and for the utility price of the dirty goods, $\zeta_d = (C_d/C)^{-1/\sigma}/C$. We show in Appendix OA2 that $\frac{\partial \dot{S}}{\partial E}\lambda_S$ is a constant equal to $\Phi \equiv \phi_L \gamma/\rho + \phi_D \gamma/(\rho+\delta)$, while ζ_d changes over time. As a result, the emissions cost τ_d^s decreases in the clean technology stock Q_c, and increases in the dirty technology stock Q_d, production labor L, and the social cost of carbon per unit of dirty good $a_d\Phi$. Marginal emission costs can also be expressed in real consumption terms (C, rather than Y_d), denoted τ_E^s. The two expressions of emission costs are related as follows: $\tau_E^s = \tau_d^s C \zeta_d / a_d = \Phi C$. This shows that the macro-economic social cost of emissions grows with consumption.

Comparing (34) to (13), we see that the market under-provides goods because of monopolistic markup α. Also, the market provides too much dirty goods because it does not price carbon emission. While the planner adjusts output in both sectors upward to correct the underproduction caused by monopoly power, she also recognizes that the production of dirty goods causes emission and increases climate damage. The latter leads to a downward adjustment of dirty production.

For the dynamic optimum, the planner's shadow value for the technology stock q_{ji}, denoted by λ_j^s, is larger than the private market's shadow value due to monopoly pricing and knowledge spillovers. The market puts too little value on innovation.

Let θ_c continue to be the relative technology of the clean sector (see (20)), and L continue to denote the total production labor. Further, similar to m_j in the decentralized equilibrium, denote by

$$m_j^s \equiv \frac{Q_j \lambda_j^s}{Q_c \lambda_c^s + Q_d \lambda_d^s}, \tag{35}$$

the share of the clean sector in the total market valuation in the planner's solution. The dynamics of the planner's solution can be characterized in terms of variables θ_c, Q_d, L, and m_c^s. Whereas in the unregulated market economy only θ_c, L, and m_c matter, in the optimal dynamics Q_d matters in addition since the dirty technology stock drives costly climate damages. The dynamics of the four variables is provided in Appendix OA2. We indicate optimal (asymptotic) steady-state values by the superscript "$**$". The next proposition characterizes the long run social optimum.

Proposition 5. *Suppose $\sigma > 1$. The welfare maximizing path converges to an asymptotic steady state with innovation in the clean sector only, $s_d = 0$, with $m_c^s \to 1$, $\theta_c \to 1$, and $L \to \rho/\mu$; the steady state dirty technology stock $Q_d^{**} \geq Q_{d,0}$ and asymptotic steady state emission cost $\tau_d^{s**} \in [0,1]$ both depend on initial conditions, where the latter is a positive function of the former as implicitly defined by*

$$\frac{\tau_d^{s**}}{(1-\tau_d^{s**})^{1/(1-\alpha)}} = \left(\frac{\alpha^{\alpha/(1-\alpha)}(1-\alpha) a_d \Phi \rho}{\mu} \right) Q_d^{**}. \tag{36}$$

Not surprisingly, welfare maximization requires a transition to a clean steady state. Because production in the dirty sector increases the atmospheric carbon stock, the climate costs in a growing economy with a non-vanishing dirty sector become infinitely large. As long as the clean sector has relatively low quality, substituting clean for dirty consumption entails a short-run cost, but in the long run this cost will always be dwarfed by the climate cost of a dirty-innovation-based growth path.[18]

[18] Technically, a dirty steady state violates the transversality condition for the carbon stock.

Since in the long run the clean sector provides all goods in the economy and all growth stems from this sector, the productivity level in the dirty goods sector becomes asymptotically irrelevant. Depending on the starting level of the dirty asset stock and on whether in the short run some dirty innovation takes place, different (but bounded) levels of dirty technology stock are compatible with a clean steady state, i.e. there is hysteresis with repect to dirty assets and emission cost. This implies stranded assets: ultimately there is no role for dirty technology in the optimal steady state, so that all initial investments become stranded in the long run.

5.2 Self-fulfilling prophecies under a Pigouvian carbon tax

The optimal transition path is characterized by a sequence of decisions with regards to which sector should be active in research. Conditional on this sequence of choices concerning the research-active sector, the market economy only faces three standard types of inefficiency – monopoly pricing, knowledge spillover, and climate externality – and the social optimum can be implemented in the market economy through regulation that directly address the three externalities.

Proposition 6. *Suppose $\sigma > 1$ and the regulated market economy innovates in the sector that the planner would choose given state variables. The social optimum can then be implemented by the simultaneous use of the following policies:*

1. *an optimal industry policy consisting of a revenue subsidy $\tau_\alpha = \alpha^{-1} - 1$ and a sector-specific technology subsidy $\tau_{qj} = s_j w / Q_j$ (per unit of q_{ji}) for the intermediate goods monopolists, and*
2. *an optimal climate policy consisting of a Pigouvian carbon tax τ_E per unit of emissions, which is set equal to the macro-economic social cost of emissions, τ_E^s, at all times.*

The above proposition shows that when the coordination failures concerning the choice of the research-active sector are somehow solved, the social optimum can be implemented in the market economy by industry and climate policy exploiting the usual three instruments. However, the next proposition shows that given such policies, the optimal steady state is not the only steady state.

Proposition 7. *Suppose $\sigma > 1$. Under the industry and climate policies in Prop 6:*

Parameter	ρ	α	μ	σ	$\theta_{c,0}$	$Q_{d,0}$	a_d	γ	ϕ_L	ϕ_D	δ	S_0	$\bar{S}$
Value	0.01	1/3	0.08	1.5	0.177	24k	0.198	0.0002	0.2	0.32	0.002	877	581

Table 1: Parameter values

1. *the regulated market economy converges to an asymptotic steady state with innovation in the clean sector only,* $s_d = 0$*, with* $m_c \to 1$*,* $\theta_c \to 1$*, and* $L \to \rho/\mu$*; (36) holds; both the steady-state dirty technology stock* $Q_d^* \geq Q_{d,0}$ *and asymptotic steady-state emission cost* $\tau_d^* \in [0,1]$ *depend on initial conditions and possibly on expectations.*
2. *given* $(Q_{c,0}, Q_{d,0})$*, there exist* $\underline{Q}_c(Q_{d,0}) > 0$ *and* $\bar{Q}_c(Q_{d,0}) \geq \underline{Q}_c(Q_{d,0})$*, and if the initial clean technology level is*
   - *(a) high such that* $Q_{c,0} \geq \bar{Q}_c(Q_{d,0})$*, equilibrium is unique and identical to the social optimum; along the transition, innovation occurs only in the clean sector;*
   - *(b) medium such that* $Q_{c,0} \in [\underline{Q}_c(Q_{d,0}), \bar{Q}_c(Q_{d,0}))$*, multiple equilibria exist; there exist a path where research is only active in the clean sector and at least one other path where research is active temporarily only in the dirty sector before switching permanently to the clean sector;*
   - *(c) low such that* $Q_{c,0} < \underline{Q}_c(Q_{d,0})$*, all paths must include temporary dirty-only research, before research switches to the clean sector permanently.*

Proposition 7 shows that, given gross substitutability between the clean and dirty goods, there exists a region of initial conditions for which multiple equilibria are possible under a Pigouvian carbon tax.[19] To provide a graphical illustration of the multiple equilibria and their welfare ranking, we consider a numerical example. We use the same parameter values for ρ, α and μ as in Section 4 and set $\sigma = 1.5$. We set $\theta_{c,0}$ to the 2019 renewable energy share in global final energy consumption (17.7 percent, IEA, 2022). Using this and the 2019 world GDP per capita ($\$11,019$ constant 2015 USD, World Bank, 2022b), we calibrate the initial dirty output $Y_{d,0}$ to $\$6,143$. Assuming a 5 percent research labor share (i.e. $L_0 = 0.95$), the initial dirty

[19] While without policy $\sigma > 2$ is required for self-fulfilling prophecies, in the regulated economy $\sigma > 1$ suffices as the industry policy in Prop 6 internalizes the social value of innovation. The overlap also becomes larger when optimal industry policy is applied. See Fig 9 in Appendix OA2.

technology stock $Q_{d,0}$ is calibrated to \$23.57k. Using $Y_{d,0}$ and the 2019 per capita carbon emission of 1.22 metric tons (World Bank, 2022a), the emission intensity a_d is calibrated to 0.198 tonnes of carbon per thousand USD. The population size is assumed to be constant at the 2019 level. For the damage function and the carbon dynamics we follow the calibration of Golosov et al. (2014), where the values of ϕ_L, ϕ_D, and δ are adjusted to annual frequency, and S_0 is updated to account for carbon emissions of recent years. An overview of the parameters is provided in Table 1.

Using the above parameter values, we solve the model numerically to find all equilibrium paths for a given initial condition $(Q_{d,0}, \theta_{c,0}, S_0)$, and compare their welfare. Using the same $(Q_{d,0}, S_0)$ values as in Table 1 while varying $\theta_{c,0}$, Figure 4 illustrates the welfare comparison of all equilibrium paths for each $\theta_{c,0}$. In particular, the red curve represents the welfare for a clean-only equilibrium path, the dashed magenta line represent an equilibrium path with temporary dirty research, and the dotted blue line represents another equilibrium path with even more dirty research. If we define a transition as the switch towards clean innovation (Lemoine, 2024), these represent an immediate transition, a short delay in transition, and a long delay in transition, respectively. We see that the entire θ_c range can be divided into three regions as stated in Proposition 7, where multiple equilibria exist in the mid range of $\theta_{c,0}$.

Intuitively, if the economy starts with enough clean capital, society reaches highest welfare when all future innovation is concentrated in the clean sector. Then complementarities in the clean sector and the need to make a transition to a clean steady state benefit clean innovation over dirty innovation from the start. If, however, the economy starts with relatively little clean assets, complementarities make innovation in the dirty sector relatively attractive in the short run, even though in the long run the economy will transition to a clean economy and assets in the dirty sector will be ultimately stranded. In this case, the cost of more asset-stranding in the long run is outweighed by the short-run benefit of larger consumption, and society is better off if research is temporarily allocated to the dirty sector.

Comparing Figures 4 and 5, we see that with high emission intensity, society prefers the clean-only path whenever multiplicity arises (Figure 4), while with low emission intensity and the clean sector starting small (Figure 5), pollution costs are

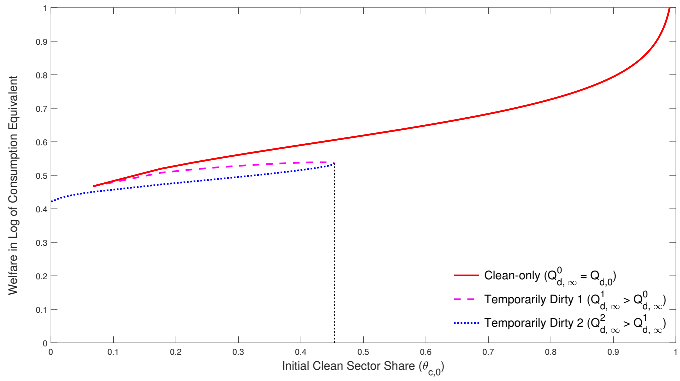

Figure 4: Welfare comparison of different equilibrium paths, calibrated emission intensity ($a_d = 0.198$)

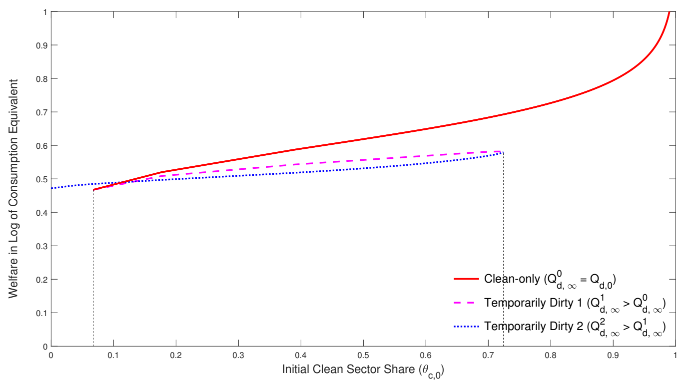

Figure 5: Welfare comparison of different equilibrium paths, counterfactual low pollution intensity ($a_d = 0.0198$, one tenth of calibrated value).

dominated by the investment complementarities and the planner prefers to delay innovation in clean. Two extreme cases help understand this result. First, in the extreme case of negligible pollution damages ($a_d \approx 0$) and $\theta_{c0} < 0.5$, investing in clean implies foregoing returns from investment complementarities and gives lower welfare than investing in dirty. Second, in the polar extreme case of infinitely large pollution damages (a_d very large), even with negligible initial clean sector clean investment dominates dirty investment. The general case lies in between these extremes, where pollution damages and intensity, as well as initial clean sector size tilt the optimum towards clean-only.

Notice from Figure 4 that the initial share suggested by the data, $\theta_{c,0} = 0.177$, falls within the region of multiple equilibria. Indeed, we find three different equilibrium paths corresponding to different levels of long-run dirty technology stock, as illustrated in Figure 6. Compared to the equilibrium path with clean research only, a path with temporary dirty research will take much longer to reach the steady state. Investing in the dirty capital thus leads to a delayed transition, and the more dirty capital the economy accumulates, the longer is the delay and the larger the stock of stranded assets.[20]

Figure 6 shows time paths of the key investment variables and the Pigouvian emission tax for the equilibria with and without delay. All equilibrium paths converge to a clean economy with the same consumption growth rate, which is also the growth rate of the long-run Pigouvian tax. In the short run, in contrast, the tax grows at different rate in the three equilibria, growing faster if there is more investment in dirty, i.e. more delay in clean investment. The equilibria with delay exploit investment complementarities in the dirty sector, which is bigger than the clean sector. Hence returns to investment are high and growth is fast, causing the social cost of carbon to grow fast, relative to the no-delay equilibrium.

To understand why two equilibria with delayed transition exist – one with a long period of dirty innovation and fast growth, and one with a short innovation period and slow growth – consider the following two opposite forces: complementarities versus pollution tax increase. With a short period before transition, polluting technology investors face only a short period with their peers investing in the same market, which comes with complementarities; but the moderate rise in the future emission tax will safeguard substantial market share even after investment complementarities stall. With a long period before transition, the opposite happens: there are prolonged short-run gains from complementarities, but dirty technologies are more quickly phased out in the longer run due to the higher tax. The result is that in both cases, short and long delay alike, the return on dirty innovation dominates that on clean innovation.

From Proposition 7 and the numerical example, we see that although taxing

[20]The finding that the steady-state polluting-sector knowledge stock depends on initial conditions is related to findings in Steger and Trimborn (2024). They discuss how in a model that is complementary to ours – a DICE climate model without endogenous technical change and evaluated in the social optimum – the steady-state stock of carbon depends on initial conditions.

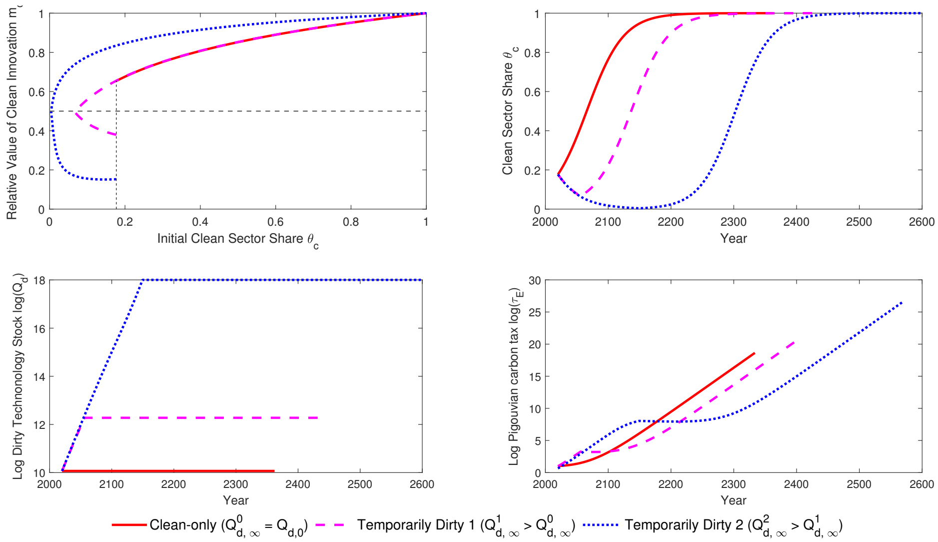

Figure 6: Fast vs delayed transitions

pollution proportional to future marginal damages effectively rules out the dirty steady state, it does not pin down a unique transition path towards a clean steady state. Whether there will be a fast transition or rather a delayed transition with a larger than optimal amount of stranded asset is again subject to coordination, and again, self-fulfilling prophecies may arise.

5.3 Coordination device

Although carbon taxation affects the relative benefit of clean and dirty innovation, the innovation coordination issue caused by strategic complementarity is ultimately a separate issue from climate externality. A carbon tax thus cannot be expected to optimally deal with both. Instead, a coordination device becomes necessary. Below we discuss policies that may serve as a coordination device, considering the case in which coordination on clean is optimal.

Minimum clean revenue guarantee Suppose the government, additionally to the policy mix of Proposition 6, guarantees the clean sector a minimum sectoral revenue of $\underline{R}$ by promising to buy clean goods if the sectoral revenue in the market falls short of the guaranteed level.[21] Since profits of intermediate goods pro-

[21]This can be interpreted as a conditional government procurement policy, which is only acti-

ducers are proportional to revenue of their sector, this sets a floor on profits for clean firms, $\pi_{ci} = \max\{(1-\alpha)(q_{ci}/Q_c)P_cY_c, (1-\alpha)(q_{ci}/Q_c)\underline{R}\}$, as well as on their shadow value of quality improvement, $\lambda_c \geq \underline{\lambda}$, and thus fully controls relative profitability of clean research. A revenue guarantee at a sufficiently high level ensures that innovation is more profitable in the clean sector, $\mu\lambda_cQ_c > \mu\lambda_dQ_d$, so that in equilibrium innovation occurs only in the clean sector. If the government can commit to a sufficiently high level of guaranteed revenue until the economy has moved out of the multiple-equilibria region, it can rule out dirty innovation at all time and coordinate investors' beliefs on the fast transition path.

Clearly, if the minimum revenue guarantee is set too high, it over-incentivizes clean innovation and distorts labor allocation. The optimal level of guaranteed revenue is the level that planner would choose, that is, $\underline{R} = P_c^*Y_c^*$. Since the policy mix of Proposition 6 ensures that the first-best allocation is a decentralized equilibrium, agents have no incentive to deviate from the first-best allocation: clean firms innovate and their equilibrium revenue exactly matches the guaranteed level. Consequently, no government payout will be necessary.

Emission cap The counterpart of a minimum revenue guarantee to the clean sector is a revenue cap to the dirty sector. When emissions are proportional to dirty output as in our model, such a revenue cap can be implemented through an emission cap. Suppose the government sets an emission cap $\bar{E}$ in addition to an industry policy as laid out in Proposition 6. Profit of a dirty intermediate goods firm is given by $\pi_{di} = (1-\alpha)(q_{di}/Q_d)(1-\tau_d)P_dY_d$, where τ_dP_d is the permit cost per unit of dirty output. Sectoral revenue net of permit costs, $(1-\tau_d)P_dY_d$, can be shown to increase with the emission cap $\bar{E}$. An emission cap thus lowers dirty firms' profit and the shadow value of dirty quality improvement λ_d, and fully controls relative profitability of clean research. A sufficiently low emission cap ensures that innovation is less profitable in the dirty sector, $\mu\lambda_dQ_d < \mu\lambda_cQ_c$, so that in equilibrium innovation only occurs in the clean sector. If the government can commit to a sufficiently low emission cap until the economy has moved out of the multiple-equilibria region, it can coordinate investors' beliefs on clean innovation and the fast transition path. To achieve the social optimum, the emission cap should be set to the emission level

vated if the market coordinates on the wrong equilibrium.

that the planner would choose at each point in time.

It may seem surprising that an emission cap can coordinate investor beliefs while a Pigouvian carbon tax cannot. The reason is that due to the Leontief production structure (fixed emissions per unit of output), an emission cap fixes the quantity of output in the dirty sector and fully controls relative profitability. It thus mutes the investment complementarity: investment in higher productivity no longer pays if the quantity policy is binding. A Pigouvian carbon tax, on the other hand, affects the marginal cost of emission and firms' marginal profit, but does not fully control the relative profitability of clean research since marginal profits still rise with total innovation in the sector – the investment complementarity may dominate the pollution tax. Although the two policies are equivalent conditional on a selected path, the equivalence breaks down when multiple equilibria with different emission levels exist.

It should be also pointed out that an emission cap only works if it indeed fully controls relative profitability. This is the case in our model where emission is proportional to dirty output. If this positive relation is broken down, such as in the case of endogenous abatement technology, an emission cap may no longer work and instead, a direct revenue cap to the dirty sector may be more appropriate.

R&D tax/subsidy A policy that affects relative profitability of clean innovation more directly is a dirty R&D tax or a clean R&D subsidy. With a dirty R&D tax, the marginal cost of dirty research becomes $(1+\tau_{rd,d})w$, which pushes the innovation regime border from $m_c = 1/2$ down to $m_c = 1/(2+\tau_{rd,d})$. Similarly, a clean R&D subsidy lowers the marginal cost of clean research $(1-\tau_{rd,c})w$ and lowers the innovation regime border to $m_c = 1/(1+(1-\tau_{rd,c})^{-1})$. Both policies lower the unique-equilibrium threshold $\bar{Q}_c$ (see proof in Appendix A). The next proposition shows that a high enough dirty R&D tax or clean R&D subsidy can lower the unique-equilibrium threshold to below the initial clean technology level.

Proposition 8. *Suppose $\sigma > 1$. Given $(Q_{c,0}, Q_{d,0})$ and the policy mix specified in Prop 6, with a sufficiently high dirty R&D tax (or clean R&D subsidy) the initial clean technology level is above the unique-equilibrium threshold ($Q_{c,0} > \bar{Q}_c$). The minimum tax rate $\tilde{\tau}_{rd,d}$ (or minimum subsidy rate $\tilde{\tau}_{rd,c}$) decreases with $Q_{c,0}$.*

If the government can commit to a sufficiently high dirty R&D tax (or clean

R&D subsidy) until the economy has moved out of the multiple-equilibria region, it can discourage dirty innovation completely and coordinate investors' beliefs on the fast transition. A dirty R&D tax has no fiscal impact and does not affect labor allocation between production and research, as no firms are actually taxed in equilibrium. In contrast, a clean R&D subsidy distorts labor allocation towards innovation, as it lowers the marginal cost of clean innovation relative to production. A clean R&D subsidy is thus an inferior coordination device as compared to a minimum clean revenue guarantee, an emission cap, or a dirty R&D tax.

Contingent high carbon tax While previous policies are unconditional policies, a policy that kicks in only when the economy innovates in the dirty sector can also serve as a coordination device. If the government can commit to a very high carbon tax contingent on dirty innovation until the economy has moved out of the mutliple-equilibria region, this eliminates dirty innovation and coordinates investor beliefs on the fast transition path. Importantly, such a contingent carbon tax must be much higher than the marginal damage of emission, that is, it must violate the Pigouvian principle, so as to eliminate undesirable equilibrium paths.

5.4 The role of commitment

A crucial condition for any coordination policy to work is policy commitment. As long as the clean market share remains in the multiple-equilibria region, delayed transition remains possible. To rule out delayed transition, investors must believe that the coordination policy will be in place as long as the economy is still in the multiple-equilibria region. This requires that the government can commit to maintaining the coordination policy over a sufficiently long horizon. This is similar in Biais and Landier (2022), who find in the context of a two-stage sequential technology spillover model that an emission cap cannot solve the coordination problem if the regulator cannot commit.

It is worth reiterating that commitment does not solve the coordination failure if the government commits to an emission tax time path that follows the time path of the first-best optimal emission tax and enacts no further coordination policy. Under commitment the tax path cannot change and private investors face the same taxes irrespective of whether the market delays clean innovation or not. Still the complementarity in investment decisions is preserved and can still dominate the effect of the tax.

6 Conclusions

We have shown that self-fulfilling prophecies arise in a standard dynamic model of DTC, according to which the equilibrium direction of innovation is determined by expectations about the future path of innovation. If goods of different sectors are sufficiently good substitutes, a demand externality causes the investments in innovation by monopolistic suppliers within a sector to be strategic complements. Applying this finding to innovation directed towards clean versus polluting technologies, we have shown that the expectation of a delay in innovation in clean technologies provides the incentive for individual investors to delay their own innovation in this direction as well, thus justifying the expectation. An equilibrium without delay in clean innovation is also possible but requires that investors do not expect delay.

Our coordination failure model provides an explanation for persistence of innovation in pollution-intensive sectors and for the slow, and possibly non-monotonic, transition to a carbon-free economy. Our welfare analysis shows excessive asset stranding in the equilibria with delay, the cost of which would have been avoided in an equilibrium with faster transition.

The implications for environmental policy are striking. With a Pigouvian pollution tax that internalizes all current and future damages, inefficient delay and asset stranding still can occur even if all agents believe that only clean technology will be used in the long run. To avoid transition delays, a coordination device is necessary. Effective coordination requires commitment to a policy that fully controls the relative profitability of innovation in the two sectors. Efficient coordination requires that the policy avoids additional distortions on other markets, such as wedges between wage cost in production versus innovation.

While our analysis focuses on the analytics of the dynamic coordination failure, the qualitative results can be expected to carry over to larger models that are more suitable for calibration and detailed policy analysis. For example, allowing for multiple sectors and several abatement options or including physical capital only affect the margins along which the economy responds to shifts in innovation direction and to environmental policy, but have no direct effect on the nature of the investment complementarities in innovation – the key driving force behind our results. Model extensions along these lines can help us evaluate the quantitative implications of the coordination failures and are an important direction for future research.

We have assumed rational expectation without uncertainty and with perfect information, by only considering equilibrium paths in which expected time path exactly coincide with realized path. As is usual in rational-expectation models with multiple equilibria, the model is silent about where the expectations come from. Future work could introduce heterogeneity of (subjective) beliefs, noisy signals about outcomes and players' actions, and explicit expectation formation.

Our model features monopolistic competition of a large number of small firms. Coordination is arguably easier in the presence of large players, whose action changes market conditions considerably. The role of large players in coordination problem has been studied in the context of currency crisis (Corsetti et al., 2004). Future research could similarly introduce large players into the DTC framework and study more targeted policies in the green transition.

References

Acemoglu, D. (2002). Directed technical change. *Review of Economic Studies 69*(4), 781–809.

Acemoglu, D., P. Aghion, L. Barrage, and D. Hémous (2023). Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution. Technical report, National Bureau of Economic Research.

Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012). The Environment and Directed Technical Change. *American Economic Review 102*(1), 131–166.

BEIS (2021, mar). Third of UK's biggest companies commit to net zero. *Press release of UK Department for Business, Energy & Industrial Strategy,*.

Benhabib, J. and R. E. Farmer (1994). Indeterminacy and increasing returns.

Benhabib, J. and R. Perli (1994). Uniqueness and Indeterminacy: On the Dynamics of Endogenous Growth. *Journal of Economic Theory 63*, 113–142.

Biais, B. and A. Landier (2022). Emission Caps and Investment in Green Technologies. *Manuscript*.

Boldrin, M. and A. Rustichini (1994). Growth and Indeterminancy in Dynamic Models with Externalities. *Econometrica 62*(2), 323–342.

Bretschger, L. and A. Schaefer (2017). Dirty history versus clean expectations: Can energy policies provide momentum for growth? *European Economic Review 99*, 170–190.

Corsetti, G., A. Dasgupta, S. Morris, and H. S. Shin (2004). Does One Soros Make a Difference? A Theory of Currency Crises with Large and Small Traders. *The Review of Economic Studies 71*(1), 87–113.

Cozzi, G. (2005). Animal spirits and the composition of innovation. *European Economic Review 49*(3), 627–637.

Cozzi, G. (2007). Self-fulfilling prophecies in the quality ladders economy. *Journal of Development Economics 84*(1), 445–464.

Draghi, M. (2024). *The future of European competitiveness: A competitiveness strategy for Europe*. European Commission.

Dugoua, E. and T. D. Gerarden (2025). Induced innovation, inventors, and the energy transition. *American Economic Review: Insights 7*(1), 90–106.

Fried, S. (2018). Climate policy and innovation: A quantitative macroeconomic analysis. *American Economic Journal: Macroeconomics 10*(1), 90–118.

Gali, J. (1996). Multiple equilibria in a growth model with monopolistic competition. *Economic Theory 8*, 251–266.

Garcia-Macia, D., C.-T. Hsieh, and P. J. Klenow (2019). How Destructive is Innovation? *Econometrica 87*(5), 1507–1541.

Gil, P. M. (2013). Animal spirits and the composition of innovation in a lab-equipment R&D model with transition. *Journal of Economics 108*(1), 1–33.

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014). Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica 82*(1), 41–88.

Hart, R. (2019). To everything there is a season: Carbon pricing, research subsidies, and the transition to fossil-free energy. *Journal of the Association of Environmental and Resource Economists 6*(2), 349–389.

Hassler, J., P. Krusell, and C. Olovsson (2021). Directed technical change as a response to natural-resource scarcity. *Journal of Political Economy 129*(11).

Hémous, D. and M. Olsen (2021). Directed technical change in labor and environmental economics. *Annual Review of Economics 13*(1), 571–597.

IEA (2022). Modern Renewables. *International Energy Agency Report "SDG7: Data and Projections"*.

Kalkuhl, M., J. C. Steckel, and O. Edenhofer (2020). All or nothing: Climate policy when assets can become stranded. *Journal of Environmental Economics and Management 100*(102214).

Krugman, P. R. (1991). History Versus Expectations. *The Quarterly Journal Of Economics 106*(2), 651–667.

Lemoine, D. (2024). Innovation-led transitions in energy supply. *American Economic Journal: Macroeconomics 16*(1), 29–65.

Matsuyama, K. (1991). Increasing returns, industrialization, and indeterminacy of equilibrium. *The Quarterly Journal of Economics 106*(2), 617–650.

Nerlich, C., P. Köhler-Ulbrich, M. Andersson, C. Pasqua, L. Abraham, K. Bańkowski, T. Emambakhsh, A. Ferrando, C. Grynberg, J. Groß, et al. (2025). Investing in europe's green future: Green investment needs, outlook and obstacles to funding the gap. Technical report, ECB Occasional Paper.

Popp, D., J. Pless, I. Hascic, and N. Johnstone (2022). Innovation and entrepreneurship in the energy sector. In *The Role of Innovation and Entrepreneurship in Economic Growth*. University of Chicago Press.

Probst, B., S. Touboul, M. Glachant, and A. Dechezleprêtre (2021). Global trends in the invention and diffusion of climate change mitigation technologies. *Nature Energy 6*(11), 1077–1086.

Smulders, S. and M. d. Nooij (2003). The impact of energy conservation on technology and economic growth. *Resources and Energy Economics 25*, 59–79.

Steger, T. and T. Trimborn (2024). Economic growth and climate change: Many trajectories, many destinations. Technical report, CESifo Working Paper.

Sturm, J. (2023). How to Fix a Coordination Failure: A 'Super-Pigouvian' Approach. *MIT mimeo*.

Thorbecke, C. (2021, sep). Over 200 companies pledge net-zero emissions by 2040 as pressure on private sector mounts. Technical report.

van der Meijden, G. and S. Smulders (2017). Carbon Lock-In: The Role of Expectations. *International Economic Review 58*(4).

van der Ploeg, F. and A. Rezai (2020). The risk of policy tipping and stranded carbon assets. *Journal of Environmental Economics and Management 100*(17400), 102258.

World Bank (2022a). CO2 emissions (metric tons per capita). *World Bank Open Data (https://data.worldbank.org/indicator/EN.ATM.CO2E.PC)*.

World Bank (2022b). GDP per capita (constant 2015 US$). *World Bank Open Data (https://data.worldbank.org/indicator/NY.GDP.PCAP.KD)*.

A Proofs

A.1 Proof Proposition 1

From (29) and (28), $\dot{\theta}_c = 0$ requires $\theta_c = 1$, $\theta_c = 0$, or $\hat{Q}_c = \hat{Q}_d$; while $\dot{m}_c = 0$ requires $m_c = 1$, $m_c = 0$, or $\hat{\lambda}_c - \hat{\lambda}_d + \hat{Q}_c - \hat{Q}_d = 0$. Since $\theta_c = 1$ is only consistent with Q_c growing faster than Q_d, this candidate steady state must have $m_c = 1$, which implies $m_c > 1/2$ so by (26) this steady state is in the clean-only regime. Similarly, a candidate steady state with $\theta_c = 0$ must also have $m_c = 0$, and by (26) the economy is the dirty-only regime. A candidate steady state with $\hat{Q}_c = \hat{Q}_d$ can only be in the simultaneous regime, which requires $m_c = 1/2$. So $\dot{m}_c = 0$ can only hold if $\hat{\lambda}_c = \hat{\lambda}_d$, which by (27) means $m_c = \theta_c$. This proves the existence of the three candidate steady states.

The rest of the proposition relies on the phase diagram (Figure 1).

As per (26), the horizontal line $m_c = 1/2$ is the regime border between clean (North) and dirty innovation (South) only. On the line innovation can be simultaneously clean and dirty. As per (27), above the 45 degree line, $m_c > \theta_c$ and $\hat{\lambda}_c > \hat{\lambda}_d$. As per (29) and (28), this implies that in the North West part of the diagram both m_c and θ_c increase; in the South East part they both decline. On the 45-degree line, $m_c = \theta_c$ and by (27) we have $\dot{m}_c/\dot{\theta}_c = 1/(\sigma - 1)$. It follows that within each regime any equilibrium path can cross the 45-degree line at most once, since the path must have the same positive constant slope whenever it crosses the 45-degree line. This slope is flatter than (equal to, or steeper than) 1 if $\sigma > 2$ ($\sigma = 2$, or $\sigma \in (1,2)$). This proves the dynamics as indicated in the figure by the small arrows.

The interior steady state ($m_c = \theta_c = 0.5$) is on the 45 degree line and on the regime border. Immediately above (below) this point in the clean (dirty) only regime, m_c and θ_c increase (decrease) so that the interior steady state can be reached neither from the clean nor the dirty innovation regime. Immediately to the left (right) of the interior steady state, $m_c = 0.5$ and $\theta_c < (>) 0.5$ and per (27) and (29) θ_c declines (increases), so that the interior steady state cannot be reached from the simultaneous innovation regime. This proves the interior steady state is unstable.

Now consider the clean candidate steady state with $\theta_c = m_c = 1$. Any equilibrium path reaching this steady state must spend a non-degenerate amount of time in the clean-only regime, in which θ_c increases over time. Since an equilibrium path can only cross the 45-degree line once, tracing back in time from the clean candidate steady state any equilibrium path must remain above ($\sigma > 2$), below ($\sigma < 2$) or on ($\sigma = 2$) the 45-degree line until it hits the regime border at some market share θ_c^{CS}. This proves that a rational expectations equilibrium path goes to the clean steady state whenever the initial market share starts above this critical level, $\theta_{c,0} \geq \theta_c^{CS}$.

By symmetry there is at least one rational expectations equilibrium path to the dirty steady state whenever $\theta_{c,0} \leq \theta_c^{DS}$.

If $\sigma > 2$, any equilibrium path towards the clean candidate steady state is above the 45-degree line as long as it remains in the clean-only regime (see the red path in the upper half of Figure 1). Thus, $\theta_c^{CS} < 1/2$. Similarly, $\theta_c^{DS} > 1/2$. This proves that for any $\theta_{c,0} \in [\theta_c^{CS}, \theta_c^{DS}]$, both candidate steady states can be reached.

If $\sigma = 2$, any equilibrium path towards the clean candidate steady state lies on the 45-degree line as long as it remains in the clean-innovation regime and $\theta_c^{CS} = 1/2$. Similarly, $\theta_c^{DS} = 1/2$. The interval $[\theta_c^{CS}, \theta_c^{DS}]$ degenerates into a single point where $m_c = \theta_c = 1/2$, which is the interior candidate steady state.

Finally, if $\sigma < 2$, any equilibrium path towards the clean candidate steady state is always below the 45-degree line as long as it remains in the clean-innovation regime. Thus, $\theta_c^{CS} > 1/2$. Similarly, $\theta_c^{DS} < 1/2$. There exists no $\theta_{c,0}$ from which both candidate steady states can be reached.

A.2 Proof Lemma 1

From (14), $\hat{Q}_j = \mu s_j$, which we substitute in (29) to get (31).

Labor market equilibrium requires wage equality, given (8) and (11):

$$w = (1-\alpha)PY/L = \mu\lambda_k Q_k. \tag{A.1}$$

From (23) and (25), the arbitrage equation (12) is written as $\hat{\lambda}_j = r - \alpha(1-\alpha)\theta_j PY/(\lambda_j Q_j)$. Substituting (A.1) and (25), we write the arbitrage equation as

$$\hat{\lambda}_j = r - \alpha\mu(Lm_k/m_j)\theta_j. \tag{A.2}$$

Using (A.2) for $j = c, d$ and substituting $\theta_d = 1-\theta_c$ and $m_d = 1-m_c$, we find:

$$\hat{\lambda}_c - \hat{\lambda}_d = \alpha\mu L\left(\frac{m_k}{m_c m_d}\right)(m_c - \theta_c). \tag{A.3}$$

Substituting (A.3) and $\hat{Q}_j = \mu s_j$ into (28), we get (30).

Time differentiation (A.1) and substitution in (A.2) gives:

$$r = \alpha\mu L\theta_k + \hat{w} - \hat{Q}_k, \tag{A.4}$$

Capital market equilibrium requires equality of the return to innovation in (A.4) and the required return in (17). Eqs. (18) and (A.1) imply $\hat{P} + \hat{C} = \hat{P} + \hat{Y} = \hat{w} + \hat{L}$, which together with (A.4) and (17) give (32).

In the simultaneous regime, $\dot{m}_c = 0$ by definition, while (A.4) holds for both $k = c$ and $k = d$. Together with (28), this gives us $\hat{Q}_c - \hat{Q}_d = \hat{\lambda}_d - \hat{\lambda}_c = \alpha\mu L(\theta_c - \theta_d)$. Substituting $\hat{Q}_c = \mu s_c$ and $\hat{Q}_d = \mu(1 - L - s_c)$, we find $s_c = (1-L)/2 + \alpha L(\theta_c - 1/2)$. From requiring $0 < s_c < 1 - L$ in the simultaneous regime, (33) follows.

A.3 Proof of Proposition 2

Prop 1 gives us the steady state values of θ_c and m_c. To find the steady state value of L, we set (32) to zero and plugging in the steady state value of θ_c.

Prop 1 already proves that the interior steady state is unstable. The stability of the corner steady states can be seen from the phase diagram in Figure 2. In the left panel, the direction of change of L and θ_c shows that the dirty steady state can only be approached from Northeast of the steady state and along the transition, both L and θ_c fall monotonically. There is thus a unique path towards the dirty steady state. Any path starting above the saddlepath implies that $L = 1$ in finite time, beyond which the arbitrage condition – i.e. the equation for $\dot{L}$ in Lemma 1 – can no longer be satisfied; any path starting below the path implies that L vanishes which is inconsistent with the transversality condition. Similar argument applies to the clean steady state.

A.4 Proof of Proposition 3

We have shown the condition for the overlap in Prop 1. For its size, we claim that $\partial\theta_c^{DS}/\partial\sigma > 0$ and $\partial\theta_c^{DS}/\partial\rho < 0$ while $\partial\theta_c^{CS}/\partial\sigma < 0$ and $\partial\theta_c^{DS}/\partial\rho < 0$. We prove here the claims for θ_c^{DS}. The proof for θ_c^{CS} is analogous.

Lemma 2. *In the dirty-only regime, for any $\theta_c > 0$, $\partial L/\partial\sigma < 0$ and $\partial L/\partial\rho > 0$.*

Proof of Lemma 2. In the θ_c-L plane (see Fig 7) the slope of the projection of the dirty equilibrium path in the dirty-only regime satisfies

$$\frac{\dot{L}}{\dot{\theta}_c} = \frac{L}{\sigma - 1}\left[\frac{(1+\rho/\mu) - (1+\alpha-\alpha\theta_c)L}{\theta_c(1-\theta_c)(1-L)}\right] \equiv f^{L\theta D}(\theta_c, L, m_c, \sigma, \rho). \qquad \text{(A.5)}$$

Evaluated at the same (θ_c, L) pair, this slope decreases with σ and increases with ρ since the term in square brackets is positive along the dirty-only equilibrium path.

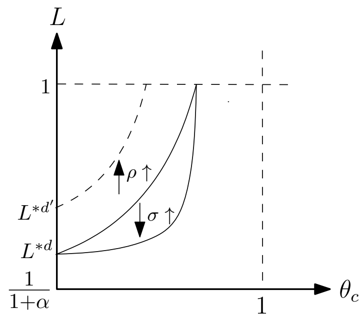

Figure 7: 2D projection in (θ_c, L)

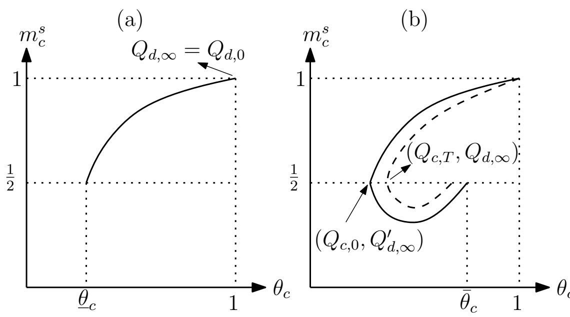

Figure 8: Thresholds for $Q_{c,0}$

Thus, equilibrium paths with different σ cross only once, namely at the steady state (since they share the same steady state) and the path becomes flatter with higher σ (see Figure 7). Equilibrium path with different ρ cannot cross, because a higher ρ corresponds to a higher steady state L so that a crossing requires the path with a higher ρ to have a flatter slope when crossing, a contradiction. Hence, the equilibrium path with a higher σ (or a lower ρ) has a lower L for any θ_c. $\square$

The projection of the equilibrium path in the dirty-only regime has the following slope in the θ_c-m_c plane:

$$\frac{\dot{m}_c}{\dot{\theta}_c} = \frac{1}{\sigma - 1}\left(\frac{m_c(1-m_c)}{\theta_c(1-\theta_c)}\right)\left[1 + \left(\frac{\alpha L}{1-L}\right)\frac{\theta_c - m_c}{m_c}\right] \equiv f^{m\theta D}(\theta_c, L, m_c, \sigma), \quad \text{(A.6)}$$

where $\theta_c > m_c$. Evaluated at the same (θ_c, m_c) pair, as long as $\sigma > 1$, this slope decreases with σ and increases with ρ since $1/(\sigma - 1)$ and L decrease with σ, while L increases with ρ (by Lemma 2). Thus again, equilibrium paths with different σ or ρ can cross only once, namely at the steady state, and the path becomes flatter as σ becomes higher or ρ becomes lower. This proves $\partial \theta_c^{DS}/\partial \sigma > 0$ and $\partial \theta_c^{DS}/\partial \rho < 0$.

A.5 Proof of Proposition 4

We prove the possible delays for a transition towards the dirty steady state. The proof for delayed transition towards the clean steady state follows similar argument.

A regime switch must occur at $m_c = 1/2$. Denote the value of θ_c and L at the final regime switch towards the dirty-only regime by θ_c^{DS} and L^{DS}. The final regime-switching point is thus given by $(\theta_c, L, m_c) = (\theta_c^{DS}, L^{DS}, 1/2)$. If (33) holds at the final regime-switching point, the economy enters the dirty-only regime from the simultaneous regime; if (33) is violated but $L^{DS} < 1$, the economy enters from the clean-only regime; and if $L^{DS} = 1$, the economy enters through stagnation.[22]

Note that θ^{DS} is the right boundary of the overlap in Prop 3 and is shown to increase with σ. We now show that L^{DS} is also increasing in σ.

Lemma 3. *$dL^{DS}(\sigma)/d\sigma > 0$.*

Proof of Lemma 3. Consider the projection of the dirty equilibrium path in the m_c-L plane within the dirty regime. From the phase diagrams, $\dot{L}/\dot{m}_c > 0$. Further, the slope of the projection is given by

$$\frac{\dot{L}}{\dot{m}_c} = \frac{L}{1-m_c}\left[\frac{\alpha\theta_c L + \rho/\mu - [(1+\alpha)L-1]}{\alpha\theta_c L - m_c\,[(1+\alpha)L-1]}\right] \equiv f^{LmD}(\theta_c, L, m_c). \tag{A.7}$$

Given (m_c, L), the slope f^{LmD} is affected by σ through θ_c:

$$\frac{\partial f^{LmD}}{\partial \theta_c} = \frac{(1-m_c)\alpha L^2}{(\dot{m}_c)^2}\left[(1-m_c)((1+\alpha)L-1) - \frac{\rho}{\mu}\right]. \tag{A.8}$$

We want to show $\partial f^{LmD}/\partial\theta_c > 0$ for all $m_c > 0$. This is equivalent to showing the equilibrium path is always above $\check{L}(m_c) \equiv (1+\frac{\rho}{\mu(1-m_c)})/(1+\alpha)$ in the m_c-L plane. From (A.7), whenever the equilibrium path and $\check{L}(m_c)$ intersect in the m_c-L plane, $f^{LmD} = L/(1-m_c) > \partial\check{L}(m_c)/\partial m_c$. Thus, the two can intersect at most once. Since they intersect at the dirty steady state, it follows that $L \geq \check{L}(m_c)$ along the equilibrium path, and from (A.8), $\partial f^{LmD}/\partial\theta_c > 0$ for $m_c > 0$.

Since given L, $\partial\theta_c/\partial\sigma > 0$ (cf Lemma 2 and Fig 7), given (m_c, L), $\partial f^{LmD}/\partial\sigma = (\partial f^{LmD}/\partial\theta_c)(\partial\theta_c/\partial\sigma) > 0$. Thus, equilibrium paths of different σ can at most intersect once in the m_c-L plane. Since they all intersect at the dirty steady state, for any $m_c > 0$, a higher σ corresponds to a larger equilibrium L. □

By Prop 3 and Lemma 3, both θ^{DS} and L^{DS} increase with σ. Both are bounded above by 1, but from (A.5) it follows $\dot{L}/\dot{\theta}_c \to +\infty$ as $\theta_c \to 1$. So $L^{DS} = 1$ can be reached without violating $\theta_c^{DS} \leq 1$. Thus, there exist $\bar{\bar{\sigma}} \in (2, \infty)$ such that $L^{DS}(\sigma) < 1$ for all $\sigma < \bar{\bar{\sigma}}$, and $L^{DS}(\sigma) = 1$ for all $\sigma \geq \bar{\bar{\sigma}}$. This proves the stagnation threshold.

[22] When $m_c = 1/2$, innovation is equally profitable in both sectors so research is active in both sectors unless 1) wage is too high and neither sector can innovate, or 2) sector-j's value of innovation m_j cannot grow as fast as in the other sector so sector j must stop innovating. The former corresponds to the case of $L^{DS} = 1$, and the latter is when (33) is violated.

At $\theta_c = \theta_c^{DS}$, the existence of the simultaneous regime requires (cf (33)):[23]

$$L \leq \Big(1+\alpha(2\theta_c^{DS}-1)\Big)^{-1} \equiv L^{RBS}(\theta_c^{DS}(\sigma)) = L^{RBS}(\sigma). \tag{A.9}$$

For $\sigma = \bar{\bar{\sigma}} > 2$, $\theta^{DS} > 1/2$ so $L^{RBS} < 1$. But $L^{DS} = 1$ so $L^{DS} > L^{RBS}$ holds. For $\sigma = 2$, $\theta_c^{DS} = 1/2$ so $L^{RBS} = 1$. But $L^{DS} < 1$ so $L^{DS} < L^{RBS}$ holds. Since $\partial L^{RBS}/\partial\sigma < 0$ and $\partial L^{DS}/\partial\sigma > 0$, there exists a unique $\bar{\sigma} \in (2, \bar{\bar{\sigma}})$ such that $L^{DS} < L^{RBS}$ for all $\sigma \in (2, \bar{\sigma})$, and $L^{DS} \geq L^{RBS}$ for all $\sigma \geq \bar{\sigma}$, proving the simultaneous-regime threshold.

If $\sigma \in (1, \bar{\sigma})$, $L^{DS} < L^{RBS}$ so that the economy reaches the final regime switch from the simultaneous regime; if $\bar{\sigma} \leq \sigma < \bar{\bar{\sigma}}$, $L^{DS} \geq L^{RBS}$ and $L^{DS} < 1$ so that the economy enters from the clean-only regime; finally, for all $\sigma \geq \bar{\bar{\sigma}}$ we have $L^{DS} = 1$ and thus there must be a stagnant period with no growth.

The next lemma shows that the simultaneous regime cannot be reached from a different regime so it is the only possible delay if $\sigma \in (1, \bar{\sigma})$, completing the proof.

Lemma 4. *If the economy is not in the simultaneous regime at time t and no unexpected shocks occur, it cannot be in the simultaneous regime after t.*

Proof of Lemma 4. Suppose at time T the economy enters the simultaneous regime from the clean-only regime. We must then have $\lim_{t\to T}\dot{m}_c < 0$ and $\lim_{t\to T} m_c = 1/2$. From (30) this requires $\theta_c > (1+\frac{1-L}{\alpha L})/2$, but this violates the second inequality in (33). The proof for the case starting with dirty only is similar. □

A.6 Proof of Proposition 5

In this proof we also refer to the dynamic equilibrium conditions provided in Appendix OA2. A steady state in the social planner's solution requires

$$\dot{\theta}_{k,\infty} = 0, \;\; \dot{L}_\infty = 0, \;\; \dot{m}^s_{k,\infty} = 0, \;\; \dot{\tau}^s_{d,\infty} = 0; \tag{A.10}$$

$$\lim_{t\to\infty} \lambda^s_j(t) q_{ji}(t) e^{-\rho t} = 0 \quad (j \in \{c,d\}); \qquad \lim_{t\to\infty} \lambda_S(t) S(t) e^{-\rho t} = 0. \tag{A.11}$$

(A.10) defines the steady state, and (A.11) is the TVCs concerning individual firm's technology stock and the aggregate carbon stock.

In the clean-only regime, $\dot{m}^s_{c,\infty} = 0$ requires $m^s_{c,\infty} = \chi_{L,c,\infty} > 1/2$. This, together with $\dot{\theta}_{c,\infty} = 0$ and $\dot{\tau}^s_{d,\infty} = 0$, requires that $\theta_{c,\infty} = 1$, which by (OA.19) leads

[23] Note that with $\theta_c^{DS} \geq 1/2$, only the second inequality in (33) is relevant.

to $\chi_{C,c,\infty} = \chi_{L,c,\infty} = 1$. $\dot{L}_\infty$ then requires $L_\infty = \rho/\mu$, while $\dot{\tau}^s_{d,\infty} = 0$ is satisfied for any $\tau^s_{d,\infty} \in [0,1]$. From (OA.20) and the steady state values $\chi_{L,c} = 1 - \chi_{L,d} = 1$ and $L = \rho/\mu$, we find (36). It is easily verified that this solution also satisfies (A.11).

In the dirty-only regime, (A.10) can only be jointly satisfied if $\theta_{c,\infty} = 0$, $\chi_{C,c,\infty} = 0$, $\chi_{L,c,\infty} = 0$, $m^s_{c,\infty} = 0$, $\tau^s_{d,\infty} = 1$, $L_\infty = \alpha\rho/\mu + 1 - \alpha$ and $s_{d,\infty} = \alpha(1-\rho/\mu)$. This solution, however, violates the TVC for the carbon stock. To see this, note that $\lambda_S S = -\lambda_{S1}S_1 - \lambda_{S2}S_2 \geq -\lambda_{S1}S_1$. In the long run, we have

$$\widehat{-\lambda_{S1}} + \hat{S}_1 - \rho = \hat{E} - \rho = \hat{Q}_d - \rho = \alpha(\mu - \rho) - \rho = \alpha\mu - (1+\alpha)\rho > 0,$$

where the inequality sign follows from the assumption of $\rho < \alpha\mu/2$ (see Prop 2). Thus, $\lim_{t\to\infty} \lambda_{S1}(t)S_1(t)e^{-\rho t} > 0$ and $\lim_{t\to\infty} \lambda_S(t)S(t)e^{-\rho t} > 0$.

In the simultaneous regime, $m^s_c = \chi_{L,c} = 1/2$ holds. From (OA.23), (OA.24) and $\dot{\chi}_{L,c} = 0$, we have $\hat{\chi}_{C,c} \propto -\hat{\tau}^s_d$. For $\dot{\chi}_{C,c} = 0$, either $\hat{\chi}_{C,c} = \hat{\tau}^s_d = 0$ or $\chi_{C,c} = 0$. If $\hat{\tau}^s_d = 0$, from $\dot{L}_\infty = 0$, (OA.29) and (OA.30) we find $\hat{Q}_{d,\infty} = 0$, which violates $s_d > 0$. If $\chi_{C,c} = 0$, together with $\chi_{L,c} = 1/2$ we obtain $\tau^s_d = 1$, which combined with (OA.19) and $\chi_{L,c} = 1/2$ requires that $\hat{Q}_c < \hat{Q}_d$. However, from (OA.24) we have $\hat{\lambda} > 0$. From $\dot{L}_\infty = 0$, (OA.29), and (OA.30) we then find $\hat{Q}_d < 0$, which also violates $s_d > 0$.

A.7 Proof of Propositions 6 and 7

Under the policies of Prop 6, all equations of motion are the same for the regulated economy and the planner's solution (see Appendix OA2). This proves Prop 6. Further, Prop 5 applies to the regulated economy. This proves part (1) of Prop 7.

The claims in part (2) of Prop 7 follow directly from the existence and definitions of the two thresholds $\underline{Q}_c$ and $\bar{Q}_c$.

Lemma 5. *Denote by $(Q_{c,T}, Q_{d,\infty})$ the technology levels when the economy enters the clean regime for the final time. $Q_{c,T}$ is a non-monotonic function of $Q_{d,\infty}$ and there exists $\bar{Q}_{c,1}$ such that $Q_{c,T} \leq \bar{Q}_{c,1}$ for all $Q_{d,\infty} \in (0,\infty)$.*

Proof of Lemma 5. From (36), $\partial\tau_{d,\infty}/\partial Q_{d,\infty} > 0$ and each $Q_{d,\infty}$ corresponds to a unique $\tau_{d,\infty}$. Given $\tau_{d,\infty}$, the steady state and the equilibrium path are uniquely determined. The equilibrium path must have a left-most intersection with the $m_c =$

$1/2$ line (see Fig 8). Denote by $\underline{\theta}_c^l$ the relative technology level at this point. Given $Q_{d,\infty}$, there is a unique $\underline{\theta}_c^l$ and thus a unique $Q_{c,T}$. Thus, $Q_{c,T}$ is a function of $Q_{d,\infty}$.

Consider $Q_{d,\infty} \to 0$. In this case, $\tau_{d,\infty} \to 0$. At the limit, this is the case without a carbon tax and $\bar{\theta}_{c,1} \in (0, 1/2)$. Denote $Q_{c,T}$ at this limit by $Q_{c,T}^0$. Clearly, $Q_{c,T}^0 \approx 0$.

Now consider $Q_{d,\infty} \to \infty$. In this case, $\tau_{d,\infty} \to 1$. By the definition of $\chi_{L,c}$, $\chi_{L,c} \to 1$ for any positive θ_c that is not close to zero. Since in the clean regime $\dot{m}_c > 0$ and $m_c > \chi_{L,c}$ must hold, $m_c \to 1$ for any positive θ_c that is not close to zero. For the equilibrium path to intersect with the regime border, $\underline{\theta}_c^l \approx 0$ must hold. Since with $Q_{c,T}^0$ and $Q_{d,\infty}$, $\theta_c \approx 0$. It follows that $Q_{c,T} \approx Q_{c,T}^0$.

Since $\bar{\theta}_{c,1} \in (0, 1/2)$ when $Q_{d,\infty} \to 0$ and $\underline{\theta}_c^l \approx 0$ when $Q_{d,\infty} \to \infty$, by continuity we can find $\tilde{Q}_{d,\infty} \in (0, \infty)$ where $\underline{\theta}_c^l \in (0, \bar{\theta}_{c,1}|_{Q_{d,\infty} \to 0})$ and thus $Q_{c,T} > 0$. Applying the same argument to all $Q_{d,\infty} \in (0, \tilde{Q}_{d,\infty})$ and $Q_{d,\infty} \in (\tilde{Q}_{d,\infty}, \infty)$, we can show that for all $Q_{d,\infty} \in (0, \infty)$, $Q_{c,T} > 0$ holds.

Since $Q_{c,T} = Q_{c,T}^0$ holds for both $Q_{d,T} \to 0$ and $Q_{d,T} \to \infty$, while $Q_{c,T} > 0$ holds for all $Q_{d,\infty} \in (0, \infty)$, $Q_{c,T}$ is non-monotonic in $Q_{d,\infty}$. Setting $\bar{Q}_{c,1} \equiv \max\{Q_{c,T} | Q_{d,\infty} \in (0, \infty)\}$ completes the proof. □

Since m_c and θ_c are slow-moving variables, the clean steady state can only be reached from the clean-only regime. Any equilibrium path is thus either a clean-only path ($Q_{d,\infty} = Q_{d,0}$) or one with temporary dirty research ($Q_{d,\infty} > Q_{d,0}$).

Threshold $\underline{Q}_c$ We check if a clean-only path is possible for an initial condition. Such a path must lie entirely within the clean-only regime. By (36), $Q_{d,\infty} = Q_{d,0}$ pins down $\tau_{d,\infty}$ and a unique equilibrium path. The projection of this path on the (θ_c, m_c) plane must have a leftmost intersection with the $m_c = 1/2$ line (see panel (a) of Fig 8). Denote θ_c and Q_c at this intersection by $\underline{\theta}_c$ and $\underline{Q}_c$, respectively. Given $Q_{d,0}$, a clean-only path is only feasible if $\theta_{c,0} \geq \underline{\theta}_c$, or equivalently, if $Q_{c,0} \geq \underline{Q}_c$.

Threshold $\bar{Q}_c$ We check if equilibrium paths with temporary dirty research are possible. Such a path must have at least one regime switch and at the final regime switch, $Q_{c,T} \geq Q_{c,0}$ must hold (= if only one regime switch).

(I) From Lemma 5, if $Q_{d,0}$ is very large such that $\partial Q_{c,T} / \partial Q_{d,\infty} < 0$, any equilibrium path with dirty research must have $Q_{c,T} < \underline{Q}_c$ (since $Q_{d,\infty} > Q_{d,0}$). Thus, if $Q_{c,0} \geq \underline{Q}_c$, clean-only path is the only equilibrium. In this case, $\bar{Q}_c = \underline{Q}_c$.

(II) Now consider the case when $Q_{d,0}$ is not too large so that $Q_{c,T}$ is non-monotonic in $Q_{d,\infty}$ for all $Q_{d,\infty} \geq Q_{d,0}$.

(IIa) If $Q_{c,0} \geq \bar{Q}_{c,1}$, by Lemma 5, no path with $Q_{c,T} \geq Q_{c,0}$ and $Q_{d,\infty} > Q_{d,0}$ can be found so there are no equilibria with temporary dirty research.

(IIb) If $Q_{c,0} < \bar{Q}_{c,1}$, by Lemma 5, at least one equilibrium path with $Q_{c,T} = Q_{c,0}$ and $Q_{d,\infty} > Q_{d,0}$ can be found. Such a path must also have a rightmost intersection $\bar{\theta}_c$ with the $m_c = 1/2$ line (see panel (b) of Fig 8). For this path to be feasible for an initial condition, $\theta_{c,0} \leq \bar{\theta}_c$ must hold. Given $Q_{d,0}$, this translates to an upper bound for $Q_{c,0}$. Consider all possible paths with one regime switch and let $\bar{Q}_{c,2}$ be the largest of all upper bounds associated with their rightmost intersection with the regime border. If $Q_{c,0} > \bar{Q}_{c,2}$, no equilibria with one regime switch exist.

(IIc) Similarly, consider all possible paths with more than one regime switches ($Q_{c,T} \in (Q_{c,0}, \bar{Q}_{c,1}]$ and $Q_{d,\infty} > Q_{d,0}$) and let $\bar{Q}_{c,3}$ be the largest of all Q_c upper bounds associated with the rightmost intersection of the paths with the regime border. If $Q_{c,0} > \bar{Q}_{c,3}$, no equilibria with more than one regime switch exist.

Summarizing (I) and (IIa)-(IIc), we conclude that no equilibria with temporary dirty research exist if $Q_{c,0} \geq \bar{Q}_c$, where $\bar{Q}_c = \underline{Q}_c$ if $Q_{d,0}$ is large (so that $\partial Q_{c,T}/\partial Q_{d,\infty} < 0$) and $\bar{Q}_c \equiv \min\{\bar{Q}_{c,1}, \max\{\bar{Q}_{c,2}, \bar{Q}_{c,3}\}\}$ otherwise.

A.8 Proof of Proposition 8

For any $Q_{d,\infty} \in (0, \infty)$, the left-most intersection of the equilibrium path and the regime border, $\underline{\theta}_c^l$, decreases as $\tau_{rd,d}$ increases and the regime border shifts down. Thus, also $Q_{c,T}$ becomes smaller. Since $Q_{c,T}$ decreases for all $Q_{d,\infty} \in (0, \infty)$, both $\underline{Q}_c$ and $\bar{Q}_{c,1}$ are decreasing in $\tau_{rd,d}$. If $\bar{Q}_c = \underline{Q}_c$, $\partial \bar{Q}_c / \partial \tau_{rd,d} < 0$ clearly holds. If $\bar{Q}_c = \min\{\bar{Q}_{c,1}, \max\{\bar{Q}_{c,2}, \bar{Q}_{c,3}\}\}$, either $\max\{\bar{Q}_{c,2}, \bar{Q}_{c,3}\}$ also decreases with $\tau_{rd,d}$ so $\partial \bar{Q}_c / \partial \tau_{rd,d} < 0$, or as $\bar{Q}_{c,1}$ falls, eventually $\bar{Q}_c = \bar{Q}_{c,1}$ holds and thus $\partial \bar{Q}_c / \partial \tau_{rd,d} < 0$.

Thus, given $Q_{d,0}$, for any $Q_c(t)$, we can find $\tilde{\tau}_{rd,d,t}$ such that $Q_c(t) = \bar{Q}_c$. As $Q_c(t)$ grows over time, $\tilde{\tau}_{rd,d}(t)$ falls. The proof for clean R&D subsidy is analogous.

OA Online Appendix for "Self-fulfilling Prophecies in the Transition to Clean Technology"

OA1 Supply side equations in the decentralized equilibrium

Substituting the firm's demand curve (9) and research technology (5), we write the Hamiltonian for the firm's maximization problem as:

$$H_{ji} = \alpha P_j L_j^{1-\alpha} q_{ji} x_{ji}^{\alpha} - P_j q_{ji} x_{ji} - w s_{ji} + \lambda_{ji} \mu Q_j s_{ji}.$$

The first order condition (FOC) for s_{ji} is (11). The FOC for x_{ji} reads

$$\alpha P_j L_j^{1-\alpha} q_{ji} x_{ji}^{\alpha-1} \alpha = P_j q_{ji},$$

which after substitution of (9) gives (10) and

$$x_{ji} = \alpha^{\frac{2}{1-\alpha}} L_j. \tag{OA.1}$$

Substituting (10) and (OA.1) into the definition of profits gives (15). Substitution of (OA.1) into the production function for Y_j gives (13).

The FOC for q_{ji} gives (12). Since $\partial \pi_{ji} / \partial q_{ji}$ is the same across firms, λ_{ji} is also the same across firms. Omitting the subscript i in (12), dividing both sides by λ_j, and substituting (15), we find

$$\dot{\lambda}_j / \lambda_j \equiv \hat{\lambda}_j = r - (1-\alpha)\alpha P_j Y_j / (Q_j \lambda_j). \tag{OA.2}$$

Using (11), (8), and (13), we find (A.2) for the research active sector. From (11) we have $\hat{\lambda}_k = \hat{w} - \hat{Q}_k$, which combined with (A.2) gives (A.4).

Finally, for any sector j, (OA.2) can be rewritten as

$$\hat{\lambda}_j = r - \alpha \mu L_k \frac{P_j L_j / \lambda_j}{P_k L_k / \lambda_k}.$$

Using $P_c Q_c = P_d Q_d$ as derived from (8) and (13) and the definition of m_c, the above equation is equivalent to

$$\hat{\lambda}_j = r - \alpha \mu L_j \frac{Q_k \lambda_k}{Q_j \lambda_j} = r - \alpha \mu L_j \frac{m_k}{m_j}. \tag{OA.3}$$

OA2 Planner's solution

From (7), we write the carbon stock S as the sum of a non-decaying stock S_1 and a decaying stock S_2, where $\dot{S}_1 = \phi_L E$ and $\dot{S}_2 = \phi_D E - \delta S_2$. To allow for symmetric expressions (across clean and dirty sectors), we write emissions as $E = a_c Y_c + a_d Y_d$ but we maintain our assumption $a_c = 0$. The current value Hamiltonian of the planner's problem is given by

$$\begin{aligned}
\mathscr{H}^{sp} &= \ln\left[\exp(-\gamma(S_1+S_2-\bar{S}))C\right] + \Omega_C\left[\left(\sum_{j\in\{c,d\}} C_j^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} - C\right] \\
&+ \sum_{j\in\{c,d\}} \Omega_{Yj}\left[L_j^{1-\alpha}\int_0^1 q_{ji}x_{ji}^{\alpha}di - Y_j\right] + \sum_{j\in\{c,d\}} \zeta_j\left[Y_j - \int_0^1 q_{ji}x_{ij}di - C_j\right] \\
&+ \sum_{j\in\{c,d\}} \int_0^1 \lambda_{ji}^s \mu s_{ji}\left(\int_0^1 q_{ji}di\right)di + \lambda_{S1}\phi_L \sum_{j\in\{c,d\}} a_j Y_j + \lambda_{S2}\left[\phi_D \sum_{j\in\{c,d\}} a_j Y_j - \delta S_2\right] \\
&+ \zeta_L\left[1 - \sum_{j\in\{c,d\}} L_j - \sum_{j\in\{c,d\}} \int_0^1 s_{ji}di\right] + \sum_{j\in\{c,d\}} \int_0^1 \xi_{s_{ji}} s_{ji}di,
\end{aligned} \tag{OA.4}$$

where C, C_j, Y_j, $\{x_{ji}\}_{i=0}^1$, L_j, $\{s_{ji}\}_{i=0}^1$ ($j \in \{c,d\}$) are the choice variables, $\{q_{ji}\}_{i=0}^1$ ($j \in \{c,d\}$) and S are the state variables, Ω_C and Ω_{Yj} are the shadow price associated with C and Y_j respectively, ζ_j and ζ_L are the shadow prices associated with the market clearing conditions, λ_{ji}^s, λ_{S1}, and λ_{S2} are the co-state variables, and finally, $\xi_{s_{ji}}$ are the shadow prices associated with the non-negativity constraints.

The FOCs are given by

$$\frac{\partial \mathscr{H}^{sp}}{\partial C}: \quad C^{-1} = \Omega_C, \tag{OA.5}$$

$$\frac{\partial \mathscr{H}^{sp}}{\partial C_j}: \quad \zeta_j = \Omega_C\left(C_j/C\right)^{-1/\sigma}, \tag{OA.6}$$

$$\frac{\partial \mathscr{H}^{sp}}{\partial x_{ji}}: \quad \zeta_j q_{ji} = \Omega_{Yj}\alpha L_j^{1-\alpha} q_{ji} x_{ji}^{\alpha-1}, \tag{OA.7}$$

$$\frac{\partial \mathscr{H}^{sp}}{\partial Y_j}: \quad \Omega_{Yj} = \zeta_j + a_j\left[\phi_L\lambda_{S1} + \phi_D\lambda_{S2}\right], \tag{OA.8}$$

$$\frac{\partial \mathscr{H}^{sp}}{\partial L_j}: \quad \zeta_L = \Omega_{Yj}(1-\alpha)Y_j/L_j, \tag{OA.9}$$

$$\frac{\partial \mathscr{H}^{sp}}{\partial s_{ji}}: \quad \zeta_L = \lambda^s_{ji}\mu Q_j + \xi_{s_{ji}}, \tag{OA.10}$$

$$\frac{\partial \mathscr{H}^{sp}}{\partial q_{ji}}: \quad \dot{\lambda}^s_{ji} = \rho\lambda^s_{ji} - \mu\int_0^1 \lambda^s_{ji}s_{ji}di - \Omega_{Y\,j}L_j^{1-\alpha}x_{ji}^{\alpha} + \zeta_j x_{ji}, \tag{OA.11}$$

$$\frac{\partial \mathscr{H}^{sp}}{\partial S_1}: \quad \dot{\lambda}_{S1} = \gamma + \rho\lambda_{S1}, \tag{OA.12}$$

$$\frac{\partial \mathscr{H}^{sp}}{\partial S_2}: \quad \dot{\lambda}_{S2} = \gamma + (\rho+\delta)\lambda_{S2}. \tag{OA.13}$$

From (OA.7) we conclude that $x_{ji} = x_j$ for all i. From (OA.10) we conclude that all producers i in sector j that are active in R&D have the same shadow price λ^s_{ji} denoted λ^s_j. Hence $\int_0^1 \lambda^s_{ji}s_{ji}di = \lambda^s_j s_j$, where s_j is aggregate R&D labor as above. Using this and $x_{ji} = x_j$ in (OA.11), we conclude that we can drop all i subscripts:

$$x_{ji} = x_j,\ s_{ji} = s_j,\ \lambda^s_{ji} = \lambda^s_j.$$

OA2.1 Social cost of carbon

Solving (OA.12) and (OA.13) we find $\lambda_{S1} = -\gamma/\rho$ and $\lambda_{S2} = -\gamma/(\rho+\delta)$, respectively. Hence, the shadow values of the two carbon stocks are constant (because of the logarithmic exponential structure as in Golosov et al. (2014)) and negative (because excess carbon causes climate damage and reduces welfare). We use Φ to denote the social cost of carbon emissions (in utility terms):

$$\Phi \equiv \lambda_S\frac{\partial \dot{S}}{\partial E} = -\lambda_{S,1}\frac{\partial \dot{S}_1}{\partial E} - \lambda_{S,2}\frac{\partial \dot{S}_2}{\partial E} = \phi_L(-\lambda_{S1}) + \phi_D(-\lambda_{S2}) = \gamma(\phi_L/\rho + \phi_D/(\rho+\delta)).$$

We define the social cost of emissions in sector j, in terms of j-goods, as

$$\tau^s_j \equiv a_j\left[\phi_L(-\lambda_{S1}) + \phi_D(-\lambda_{S2})\right]/\zeta_j.$$

Thus, we find that in the optimum the emission costs equal:

$$\tau^s_d = a_d\Phi/\zeta_d; \quad \tau^s_c = a_c\Phi/\zeta_c = 0, \tag{OA.14}$$

where we refer to τ^s_j as the tax and introduce the zero tax in the clean sector τ^s_c to allow symmetry in our expressions below.

OA2.2 Optimal input mix and static allocation

Because $x_{ji} = x_j$, the production function and goods market equilibrium can be written as, respectively, $Y_j = L_j^{1-\alpha}Q_j x_j^{\alpha} = C_j + Q_j x_j$. Substituting (OA.7) and (OA.8), we find expressions (34) for the production function and we find the consumption-output ratio C_j/Y_j as a function of the tax:

$$Y_j = Q_j L_j[\alpha(1-\tau_j^s)]^{\alpha/(1-\alpha)} \tag{OA.15}$$

$$C_j = Y_j[1-\alpha(1-\tau_j^s)]. \tag{OA.16}$$

Let $\chi_{C,j}$ denote the share of goods j in total value of consumption and χ_{Lj} denote the share of production labor hired in sector j, that is

$$\chi_{C,j} \equiv \frac{\zeta_j C_j}{\Omega_C C}, \qquad \chi_{L,j} \equiv \frac{L_j}{L}. \tag{OA.17}$$

$\chi_{C,j}$ is thus the direct counterpart of the expenditure share in the decentralized equilibrium and $\chi_{L,j}$ the production labor share. From (OA.9), (OA.5), (OA.16), and (OA.17), we express the shadow price of labor as:

$$\zeta_L = \left(\frac{(1-\tau_j^s)(1-\alpha)}{1-\alpha(1-\tau_j^s)}\right)\frac{\chi_{Cj}}{L_j} \tag{OA.18}$$

We write the six equations (OA.9), (OA.7), (OA.6), (OA.8), (OA.15), and (OA.16) in relative terms (clean versus dirty) and solve for $\zeta_r, \Omega_{Yr}, C_r, Y_r, L_r, x_r$ in terms of τ_j^s and Q_r. Using these solutions with (20), the definition of θ_c, we find

$$\frac{\chi_{C,c}}{\chi_{C,d}} = \frac{\theta_c}{\theta_d}(1-\tau_d^s)^{-\frac{\sigma-1}{1-\alpha}}, \qquad \frac{\chi_{L,c}}{\chi_{L,d}} = \frac{\theta_c}{\theta_d}(1-\tau_d^s)^{-\frac{\sigma-\alpha}{1-\alpha}}\frac{1-(1-\tau_d^s)\alpha}{1-\alpha}. \tag{OA.19}$$

While in the decentralized equilibrium the relative technology fully captures the economic incentive for clean production and consumption, in the planner's solution these economic incentives must be augmented by the technology's contribution to carbon emission. Compared to the decentralized equilibrium, with the same level of relative technology θ_c, the planner will allocate more labor to the clean sector and consume a larger share of clean goods.

The allocation of labor for R&D is governed by (OA.10). For the research active sector, $\zeta_L = \lambda_k^s \mu Q_k$ must hold, while $\zeta_L > \lambda_{-k}^s \mu Q_{-k}$ holds for the research inactive sector. Using the definition of m_c^s, (35), we find that innovation is only active in the clean (dirty) sector if $m_c^s > 1/2$ (if $m_c^s < 1/2$). Thus, $m_c^s = 1/2$ separates the innovation regimes, just like in the decentralized equilibrium.

OA2.3 Static expression for the optimal tax

From (OA.5), (OA.6), (OA.17), and the definition of τ_d^s, we find $\tau_d^s = a_d\Phi/\zeta_d = a_d\Phi C_d/\chi_{C,d}$. Substituting $1/\chi_{C,d} = 1 + \chi_{C,c}/\chi_{C,d}$, (OA.16), (OA.15), and $L_d = \chi_{L,d}L$, we write:

$$\tau_d^s = a_d\Phi(\chi_{L,d}L)Q_d[\alpha(1-\tau_d^s)]^{\alpha/(1-\alpha)}(1+\chi_{C,c}/\chi_{C,d})[1-\alpha(1-\tau_d^s)].$$

Since, from (OA.19) we find $\chi_{C,c}/\chi_{C,d} = (\chi_{L,c}/\chi_{L,d})(1-\tau_d^s)(1-\alpha)/(1-\alpha(1-\tau_d^s))$, we can write:

$$\tau_d^s = a_d\Phi LQ_d[\alpha(1-\tau_d^s)]^{\alpha/(1-\alpha)}\left[(1-\alpha(1-\tau_d^s))\chi_{L,d} + (1-\tau_d^s)(1-\alpha)\chi_{L,c}\right]. \tag{OA.20}$$

This equation gives a relationship between the tax and other key variables. We are interested in the relationship with θ_c, L, and Q_d. We therefore substitute $\chi_{L,d} = 1-\chi_{L,c}$ at the RHS and divide both sides by τ_d^s. It can be easily seen that then the RHS declines with τ_d^s and with $\chi_{L,c}$, where the latter itself increases with τ_d^s and θ_c. Hence, there is a unique solution for τ_d^s as a function of θ_c, L, and Q_d with the following properties:

$$\tau_d^s = \tilde{\tau}(\theta_c, a_d\Phi LQ_d) \in [0,1), \quad \tilde{\tau}_1 < 0, \quad \tilde{\tau}_2 > 0. \tag{OA.21}$$

OA2.4 Dynamic allocation

Time differentiating (20) and (OA.17), we find

$$\dot{\theta}_c = (\sigma-1)\theta_c(1-\theta_c)\left(\hat{Q}_c - \hat{Q}_d\right), \tag{OA.22}$$

$$\dot{\chi}_{C,c} = (\sigma-1)\chi_{C,c}(1-\chi_{C,c})\left[\hat{Q}_c - \hat{Q}_d + \frac{1}{1-\alpha}\frac{\tau_d^s}{1-\tau_d^s}\hat{\tau}_d^s\right], \tag{OA.23}$$

$$\dot{\chi}_{L,c} = (\sigma-1)\chi_{L,c}(1-\chi_{L,c})\left[\hat{Q}_c - \hat{Q}_d + \left(\frac{1}{1-\alpha} + \frac{1}{(\sigma-1)(1-\alpha(1-\tau_d^s))}\right)\frac{\tau_d^s}{1-\tau_d^s}\hat{\tau}_d^s\right]. \tag{OA.24}$$

From (OA.10) and the definition of m_j^s, (35), we derive as the planner's counterpart of (26) that sector j is the research-active sector if its social market value exceeds cost: $m_j^s > \kappa_j \Leftrightarrow s_j > 0$. Since $j = k$ denotes the research-active sector, we must have $m_k^s \geq \kappa_k$ and $\zeta_L = \mu \lambda_k^s Q_k$. This implies:

$$k = \begin{cases} c, & \text{if } m_c^s > 1/2; \\ d, & \text{if } m_c^s < 1/2. \end{cases} \tag{OA.25}$$

$$(\hat{Q}_c, \hat{Q}_d) = \begin{cases} (\mu(1-L), 0), & \text{if } k = c; \\ (0, \mu(1-L)), & \text{if } k = d. \end{cases} \tag{OA.26}$$

Substituting (OA.7), (OA.9), and (OA.15) into (OA.11) gives $\hat{\lambda}_j^s = \rho - \mu s_j - \zeta_L L_j / \lambda_j^s Q_j$. For the research active sector ($j = k$) we have $s_k > 0$ and $\zeta_L = \lambda_k^s \mu Q_k$ from (OA.10). Together with the definition of m_j^s, (35), this gives:

$$\hat{\lambda}_j^s = \rho - \mu s_j - \frac{m_k^s}{m_j^s} \mu L_j. \tag{OA.27}$$

For $j = k$, (OA.27) implies $\hat{\lambda}_k^s = \rho - \mu(L_k + s_k)$. To derive the dynamics for m_j^s, note that $\hat{Q}_j = \mu s_j$ and $\hat{m}_k^s - \hat{m}_j^s = (\hat{\lambda}_k^s + \hat{Q}_k) - (\hat{\lambda}_j^s - \hat{Q}_j) = (m_k^s/m_j^s - L_k/L_j)\mu L_j$ where the second equality follows from (OA.27). Using $L_j = \chi_{L,j} L$ from (OA.17), we find

$$\dot{m}_c^s = \begin{cases} m_c^s (m_c^s - \chi_{L,c}) \mu L, & \text{if } k = c; \\ 0, & \text{if } k = c, d; \\ (1 - m_c^s)(m_c^s - \chi_{L,c}) \mu L, & \text{if } k = d. \end{cases} \tag{OA.28}$$

To derive the dynamics for L, we combine (OA.10) and $\hat{\lambda}_k^s = \rho - \mu(L_k + s_k)$ to arrive at $\hat{\lambda}_k^s + \hat{Q}_k = \hat{\zeta}_L = \rho - \mu L_k$, while time differentiating (OA.18) implies $\hat{\zeta}_L = \hat{\chi}_{C,c} - \hat{\chi}_{L,c} - \hat{L}$. Hence, we arrive at $\hat{L} = \mu L_k - \rho + \hat{\chi}_{C,c} - \hat{\chi}_{L,c}$. Substituting (OA.23) and (OA.24), we find:

$$\dot{L} = L \Big[\mu L_k - \rho + (\sigma - 1)(\chi_{L,c} - \chi_{C,c}) \left(\hat{Q}_c - \hat{Q}_d \right) + \left(\frac{(\sigma-1)(\chi_{L,c} - \chi_{C,c})}{1-\alpha} - \frac{1 - \chi_{L,c}}{1 - \alpha(1 - \tau_d^s)} \right) \frac{\tau_d^s}{1 - \tau_d^s} \hat{\tau}_d^s \Big]. \tag{OA.29}$$

Finally, we derive the dynamics of τ_d^s. From (OA.8) and (OA.9), we find $\zeta_d(1 - \tau_d^s) = \zeta_L L_d / ((1-\alpha) Y_d)$ and after using the definition of τ_d^s to eliminate ζ_d using

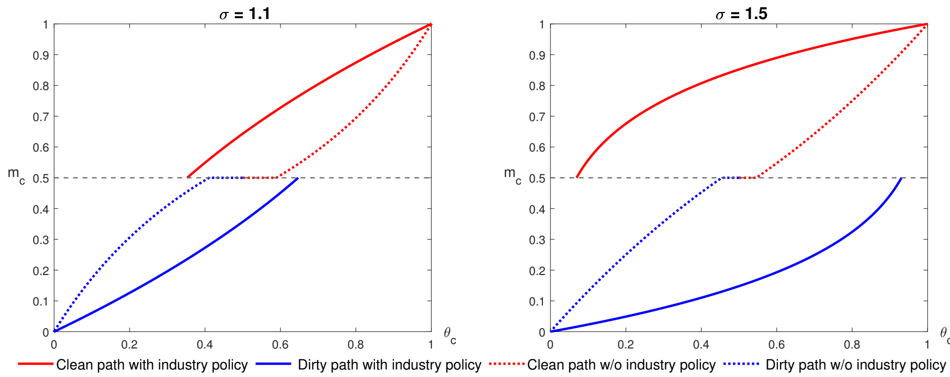

Figure 9: Size of overlap under industrial policy

(OA.10) and (OA.15) to eliminate ζ_L and Y_d, respectively, we derive

$$(1-\alpha)a_d\Phi = \tau_d^s[(1-\tau_d^s)\alpha^{\alpha}]^{-1/(1-\alpha)}\lambda_k^s\mu Q_k/Q_d.$$

Time differentiating this equation and substituting (OA.27) to eliminate $\hat{Q}_k$, we find the dynamics of the tax:

$$\dot{\tau}_d^s = -\frac{(1-\alpha)\tau_d^s(1-\tau_d^s)}{1-\alpha(1-\tau_d^s)}\left(\rho-\mu L_k-\hat{Q}_d\right). \tag{OA.30}$$

OA2.5 Characterising dynamics

Using (OA.19), (OA.21), and (OA.30) to eliminate χ_L, χ_C, τ_d^s, and $\dot{\tau}_d^s$, respectively, we find that (OA.22), (OA.26), (OA.28), and (OA.29) constitute four differential equations in four variables, namely θ_c, m_c, L, and Q_d.

The projection of the dynamics in the (θ_c, m_c^s) plane, used in Figure 8, can be characterised as follows. From (OA.25) we find the line $m_c^s = 1/2$ as the regime border which acts as the $\dot{\theta}_c = 0$ locus: above (below) the line there is clean (dirty) innovation only, and θ_c increases (decreases) over time if and only if $\sigma > 1$. Next, from (OA.28) we derive $m_c^s = \chi_{L,c}$ as the $\dot{m}_c^s = 0$ locus, with m_c increasing (decreasing) over time above (below) the locus. This locus is not a fixed line in the plane, since $\chi_{L,c}$ depends on not only θ_c but also τ_d^s, see (OA.19), which depends on the whole dynamics of the system, cf. (OA.21). Nevertheless, (OA.19) shows that $\chi_{L,c} > \theta_c$ and $\chi_{L,c} \to \theta_c$ at the corners $\theta_c \to 0$ and $\theta_c \to 1$, so that the $\dot{m}_c^s = 0$ locus cuts the 45 degree line in the corners and is above the 45 degree line for θ_c.

From (OA.22) and (OA.28), we derive that the slope of the optimum path, $\dot{m}_c^s/\dot{\theta}_c$, in the corner $(1,1)$ equals 0. Hence the optimal path must approach the clean steady state from the south west.

OA2.6 Regulated market economy

With a tax τ_E on carbon emission (in real terms), the profit of final goods producers becomes

$$\pi_j = P_j Y_j - wL_j - \int_0^1 P_{ji} x_{ji} di - \tau_E P a_j Y_j = (1-\tau_j) P_j Y_j - wL_j - \int_0^1 P_{ji} x_{ji} di, \tag{OA.31}$$

where $\tau_j = a_j \tau_E P/P_j$ is the emission tax in terms of revenue (i.e. formulated as a value-added tax). Maximizing profits subject to the production function leads to a modified factor demand:

$$w = (1-\alpha)(1-\tau_j) P_j \frac{Y_j}{L_j}, \tag{OA.32}$$

$$P_{ji} = \alpha(1-\tau_j) P_j L_j^{1-\alpha} q_{ji} x_{ji}^{\alpha-1}. \tag{OA.33}$$

With this modified factor demand and the industry policy specified in the proposition, the Hamiltonian of the intermediate goods producers becomes

$$H_{ji} = (1+\tau_\alpha)\alpha(1-\tau_j) P_j L_j^{1-\alpha} q_{ji} x_{ji}^{\alpha} - P_j q_{ji} x_{ji} - ws_{ji} + \tau_{qj} q_{ji} + \lambda_{ji} \mu Q_j s_{ji}. \tag{OA.34}$$

where τ_α is the revenue subsidy, and τ_{qj} is the sector-specific technology subsidy. Accordingly, (10)-(12) change to

$$P_{ji} = \frac{1}{\alpha}(1+\tau_\alpha)^{-1} P_j q_{ji} = P_j q_{ji}, \tag{OA.35}$$

$$\mu Q_j \lambda_{ji} \leq w \perp s_{ji} \geq 0, \tag{OA.36}$$

$$\dot{\lambda}_{ji} = r\lambda_{ji} - \tau_{qj} - \frac{\partial \pi_{ji}}{\partial q_{ji}}. \tag{OA.37}$$

Combining the above results with (16) and (17), and using the symmetry result $x_{ji} = x_j, \lambda_{ji} = \lambda_j$, we can now summarize the regulated market economy by the following equations:

$$P_j/P = (C_j/C)^{-1/\sigma} \tag{OA.38}$$

$$1 = (1-\tau_j)\alpha L_j^{1-\alpha} x_j^{\alpha-1} \alpha(1+\tau_\alpha) \tag{OA.39}$$

$$w = (1-\tau_j) P_j (1-\alpha) Y_j / L_j \tag{OA.40}$$

$$w = \lambda_k \mu Q_k \tag{OA.41}$$

$$\hat{\lambda}_j - \hat{P} - \hat{C} = \rho - s_{qj}/\lambda_j - (1-\tau_j) P_j L_j^{1-\alpha} x_j^{\alpha} (1-\alpha)\alpha(1+\tau_\alpha) \tag{OA.42}$$

Combining the optimality conditions of the social planner's problem (OA.5)-(OA.11), using symmetry $x_{ji} = x_j, \lambda^s_{ji} = \lambda_j$ and the definition $\tau^s_j \equiv a_j \left[\phi_L(-\lambda_{S1}) + \phi_D(-\lambda_{S2})\right]/\zeta_j$, we find that the social planner's solution satisfies the following equations:

$$\zeta_j C = (C_j/C)^{-1/\sigma}, \tag{OA.43}$$

$$1 = (1-\tau_j)\alpha L_j^{1-\alpha} x_j^{\alpha-1}, \tag{OA.44}$$

$$\zeta_L = (1-\tau_j)\zeta_j (1-\alpha) Y_j / L_j, \tag{OA.45}$$

$$\zeta_L = \lambda^s_k \mu Q_k, \tag{OA.46}$$

$$\hat{\lambda}^s_j = \rho - \mu s_j - (1-\tau_j)\zeta_j L_j^{1-\alpha} x_j^{\alpha}(1-\alpha). \tag{OA.47}$$

Comparing (OA.43)-(OA.47) for the optimal economy to (OA.38)-(OA.42) for the regulated economy, we find that the latter replicates the former if the tax policies of proposition are imposed, $(1+\tau_\alpha)\alpha = 1$, $\tau_{qj} = ws_j/Q_j$, and $\tau_j = \tau^s_j$. Note that this implies $P_j/PC = \zeta_j$, $w/PC = \zeta_L$, $\lambda_j/PC = \lambda^s_j$, i.e. the real market prices in utility terms (market prices divided by P to make C the unit of account and then multiplied by marginal utility $1/C$ to make utility the unit of account) equal the corresponding shadow prices.

OA3 General condition for the overlap

This appendix, first, generalizes the production and innovation technology to allow for more general complementarities in innovation and, second, relaxes the patent length assumption to allow for variable patent length.

OA3.1 Generalizing the sources of complementarity

We generalize the model in three ways to allow for multiple sources of investment complementarities. First, we allow a direct effect of intermediate firms' innovation on productivity in their sector, by generalizing the final good production to be

$$Y_j = (Q_j^{\varepsilon} L_j)^{1-\alpha} \int_0^1 q_{ji} x_{ji}^{\alpha} di, \tag{OA.48}$$

where $\varepsilon \geq 0$ measures how labor-augmenting the direct innovation spillovers are ($\varepsilon = 0$ brings us back to the main text model).

Second, we allow for a more general input-output structure by assuming intermediate goods production requires both own sector goods and general goods. The unit cost (or equivalently, its monopoly price divided by markup) of an intermediate in sector j with quality q_{ji} is now

$$q_{ji} P_j^{\omega} P^{1-\omega} = \alpha P_{ji}, \tag{OA.49}$$

where $\omega \geq 0$ measures the share of own sector inputs in the production of specific inputs ($\omega = 1$ brings us back to the main text model; Acemoglu et al (2012) choose $\omega = 0$).

Third, we allow intersectoral knowledge spillovers in innovation such that

$$\dot{q}_{ji} = \mu s_{ji} Q_j^{\eta+\chi} Q_{-j}^{\chi} (Q_c + Q_d)^{1-\eta-2\chi}, \tag{OA.50}$$

where $Q_c + Q_d$ is the general knowledge stock, χ is the degree of cross-sectoral spillovers and η denotes how much more own-sector knowledge enhances research productivity than other-sector knowledge; we have maintained the linear homogeneity that was also assumed in the main text. The model presented in the main text can then be considered a special case where $\eta = 1, \chi = 0$.

Lemma OA1. *In a static equilibrium, intermediate goods profits are linear in firms own quality q_{ji}, i.e. $\pi_{ji} = \bar{\pi}_j q_{ji}$ with*

$$\bar{\pi}_r = (Q_r)^{\psi}, \psi \equiv (1+\varepsilon)(\sigma - 1)\frac{1-\alpha}{1-\omega\alpha} - 1, \tag{OA.51}$$

where $\bar{\pi}_r \equiv \bar{\pi}_c / \bar{\pi}_d$ and $Q_r \equiv Q_c / Q_d$, while relative R&D costs are

$$\frac{ws_{di}/\dot{q}_{di}}{ws_{ci}/\dot{q}_{ci}} = \mu_r (Q_r)^{\eta}. \tag{OA.52}$$

Proof of Lemma OA1. We determine static equilibrium, i.e. the allocation of labor and profits, given the state variables q_{ji} and given the amount of labor in production L.[1]

[1]Using this static allocation, below we turn to the dynamic equilibrium to determine the alloca-

From demand for intermediates ($P_j \partial Y_j / \partial x_{ji} = P_{ji}$) and supply (OA.49) we find $x_{ji} = \alpha^{2/(1-\alpha)} Q_j^{\varepsilon} L_j (P_j/P)^{(1-\omega)/(1-\alpha)}$. Plugging this into the production function we find $Y_j = \alpha^{2\alpha/(1-\alpha)} Q_j^{\varepsilon+1} L_j (P_j/P)^{(1-\omega)\alpha/(1-\alpha)}$. Hence, in relative terms:

$$x_r = Q_r^{\varepsilon} L_r (P_r)^{(1-\omega)/(1-\alpha)},$$

$$Y_r = Q_r^{\varepsilon+1} L_r (P_r)^{(1-\omega)\alpha/(1-\alpha)}.$$

Demand for labor implies $P_j \partial Y_j / \partial L_j = P_j (1-\alpha) Y_j / L_j = w$, or in relative terms

$$L_r = P_r Y_r.$$

Demand for Y-goods implies:

$$Y_r = (P_r)^{-\sigma}.$$

Hence we have four equations in P_r, L_r, Y_r, x_r which can be solved in terms of Q_r.

$$P_r = (Q_r)^{-(\varepsilon+1)(1-\alpha)/(1-\omega\alpha)}$$
$$L_r = (Q_r)^{(\sigma-1)(\varepsilon+1)(1-\alpha)/(1-\omega\alpha)}$$
$$x_r = (Q_r)^{\varepsilon+(\varepsilon+1)[(\sigma-1)(1-\alpha)-(1-\omega)]/(1-\omega\alpha)}$$

Now we turn to profits of intermediate firms. Since the markup is $1/\alpha$, profits are $\pi_{ji} = (1-\alpha) P_{ji} x_{ji}$. and the price P_{ji} from (OA.49), we find $\pi_{ji} = [(1-\alpha)\alpha^{-1} x_{ji} P_j^{\omega} P^{1-\omega}] q_{ji} \equiv \bar{\pi}_j q_{ji}$, where the latter step uses the result that x_{ji} in equilibrium is the same across firms. This shows that profits are linear in own quality q_{ji}, which is stated in the lemma. Plugging in the solution for x_{ji} and taking relative variables, we find $\bar{\pi}_r = x_r (P_r)^{\omega}$ which together with above solutions gives (OA.51).

From (OA.50) we directly find (OA.52). □

Hence ψ reflects investment complementarities in production: if $\psi > 0$, an increase in relative knowledge stocks increases relative marginal profits (in the main text, $\omega = 1, \varepsilon = 0$ so that $\psi = \sigma - 2$). Complementarities arise from (i) demand externalities (σ) (ii) input-output mulitipliers (ω) and (iii) direct productivity spillovers (ε). Furthermore, η reflect investment complementarities in innovation: if $\eta > 0$ investment in sector j reduces the cost of subsequent investment more than in the other sector.

tion of labor over production and innovation and the resulting dynamics of the state variable.

Lemma OA2. *SFPs in the unregulated market economy require* $\psi > \max\{0, -\eta\}$.

Proof of Lemma OA2. This proof turns to the dynamics of the model and exploits the static equilibrium solutions in terms of the state variables q_{ji} from the revious proof. The intermediate good producer's investment problem of choosing s_{ji} has the following Hamiltonian:

$$H_{ji} = \bar{\pi}_j q_{ji} - w s_{ji} + \lambda_{ji} \mu \bar{Q}_j s_{ji}.$$

where $\bar{Q}_j = Q_j^{\eta+\chi} Q_{-j}^{\chi} (Q_j + Q_{-j})^{1-\eta-2\chi}$ is the productivity of research labor. The firm takes variables without i subscript as given. Optimality conditions are:

$$\bar{Q}_j \mu \lambda_{ji} \leq w \perp s_{ji} \geq 0 \tag{OA.53}$$

$$\hat{\lambda}_{ji} = r - \bar{\pi}_j / \lambda_{ji} \tag{OA.54}$$

The two conditions show that all firms within a sector have the same shadow value of quality, $\lambda_{ji} = \lambda_j$. We define two variables, z_c and m_c:

$$z_c \equiv \frac{\bar{\pi}_c \bar{Q}_c}{\bar{\pi}_c \bar{Q}_c + \bar{\pi}_d \bar{Q}_d} = \frac{(Q_r)^{\psi+\eta}}{1 + (Q_r)^{\psi+\eta}}, \tag{OA.55}$$

$$m_c \equiv \frac{\lambda_c \bar{Q}_c}{\lambda_c \bar{Q}_c + \lambda_d \bar{Q}_d} = \frac{\lambda_r (Q_r)^{\eta}}{1 + \lambda_r (Q_r)^{\eta}}. \tag{OA.56}$$

Variable z_c captures current (green) market conditions. It is a predetermined state variable, i.e. a transformation of the relative technology state variable Q_r. The transformation ensures that z_c captures all channels through which the state variable affects the return to innovation: complementarities in production (ψ) and in innovation (η). Variable m_c captures future (green) market conditions. It is a forward-looking variable constructed such that its value directly pins down which innovation is active. Clean (dirty) innovation requires future green market conditions to be sufficiently good (poor) according to:[2]

$$m_c > (<) 1/2 \Leftrightarrow \hat{Q}_r > (<) 0.$$

From optimality condition (OA.54) we derive the relative growth rates $\hat{\lambda}_r =$

[2]From (OA.53) we derive the regime border condition $\lambda_r \bar{Q}_r \mu_r > (<) 1 \Leftrightarrow \hat{Q}_r > (<) 0$ which in terms of m_c gives the expression.

$\frac{\bar{\pi}_d}{\lambda_d}\left(1-\frac{\bar{\pi}_r}{\lambda_r}\right)$ which in terms of our new variables reads:[3]

$$\hat{\lambda}_r = \left(\frac{\bar{\pi}_d/\lambda_d}{(1-z_c)m_c}\right)(m_c - z_c).$$

To derive the dynamics of the model in terms of z_c and m_c, we time differentiate (OA.55) and (OA.56):

$$\dot{z}_c = z_c(1-z_c)(\psi+\eta)\hat{Q}_r \tag{OA.57}$$

$$\dot{m}_c = m_c(1-m_c)(\hat{\lambda}_r + \eta\hat{Q}_r), \tag{OA.58}$$

We now build the phase diagram in (z_c, m_c) plane. The regime border is the horizontal line $m_c = 1/2$. We first consider $\psi + \eta < 0$ and show that this rules out SFPs. If $m_c < 1/2$, innovation is brown, Q_r declines and z_c grows. Symmetric for $m_c > 1/2$. Hence the interior steady state with simultaneous research and $m_c = z_c = 1/2$ is stable, the corner steady states can never be reached, and no SFPs can arise.

We next show that an overlap requires $\psi > 0$. Assume $\psi + \eta > 0$. The slope of any time path is given by $\dot{m}_c/\dot{z}_c$. On the 45 degree line (with $m_c = z_c$ and hence $\hat{\lambda}_r = 0$) this slope boils down to:

$$\left.\frac{\dot{m}_c}{\dot{z}_c}\right|_{m_c=z_c} = \frac{\eta}{\psi+\eta}. \tag{OA.59}$$

This means that, unless $\eta/(\psi+\eta) = 1 \iff \psi = 0$, an equilibrium path can cross the 45 degree line only once. When tracing back the equilibrium path from a corner steady state (either the dirty staedy state $m_c = z_c = 0$ or the clean one $m_c = z_c = 1$), we start on the 45 degree line and never cross again; when the slope is smaller than 1, the path from the dirty (clean) steady state crosses the regime border to the right (left) of the 45 degree line, implying an overlap. Hence the condition for SFPs is $\eta/(\psi+\eta) < 1 \Leftrightarrow \psi > 0, \psi+\eta > 0$. □

Remark. This proof only uses the investment conditions and does not need consumer intertemporal utility maximization. This is becuase we only need to solve for relative variables. When we want to solve for all variables, in particular total - rather than relative - investment, as measured by $1-L$, we need the savings block of the model.

[3]Note $\bar{\pi}_r/\lambda_r = z_r/m_r = [z_c/(1-z_c)]/[m_c/(1-m_c)]$.

OA3.2 Variable patent length

While Acemoglu et al. (2012) assume one-period patents and our main text model assumes infinite patent length, in reality patents often last between 15 and 20 years. To model elementary aspects of patent protection issues, we assume all intermediate firms face a risk of losing their profits permanently because of patent infringement.[4] The infringement event occurs at Poisson rate ι, so that the arbitrage equation (12) now contains a risk premium:

$$\dot{\lambda}_{ji} = (r+\iota)\lambda_{ji} - \frac{\partial \pi_{ji}}{\partial q_{ji}}, \tag{OA.60}$$

which implies that profits are discounted with the interest rate plus the infringement risk ι to calculate the value of investment λ_{ji}.

The patent infringement rate does not affect the analysis in Section 3. Intuitively, because patent infringement occurs with the same probability in all sectors, it does not affect the direction of investment.

However, the patent infringement rate affects the speed of overall investment as analyzed in Section 4. Following the procedure of Section A.2, we derive the counterparts of (A.2) and (A.4),

$$\hat{\lambda}_j = r + \iota - \alpha\mu L(m_k/m_j)\theta_c, \tag{OA.61}$$

$$r = \alpha\mu L\theta_k + \hat{w} - \hat{Q}_k - \iota. \tag{OA.62}$$

Continuing the same procedure as in Section A.2 to derive the reduced-form equilibrium dynamics, we find that in all equations in Lemma 1, ρ is replaced by $\rho+\iota$. Intuitively, a higher probability of loosing the patent right reduces investors' horizon as does an increase in the discount rate, so that the sum of discount rate and patent infringement rate governs the speed of innovation. The effect of a change in ρ and a change in ι are the same with respect to the equilibrium dynamics analyzed in Section 4. Hence, we conclude that a shorter average patent length (increase in ι) makes the overlap smaller.

Figure 10 shows the projection of the equilibrium paths for different time pref-

[4] This modelling assumes that infringement is exogenous and uniform across firms; firms who "steal" the patent are immediately in the same position as robbed incumbent. A full modelling would require specifying who is successful in infringement, whether this costs effort etc. Moreover, (legal) patent length is not the same as (illegal) infringement. We leave these details for further research.

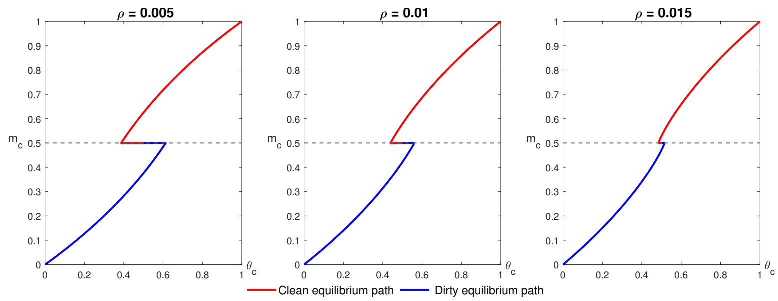

Figure 10: Equilibrium paths with different ρ values ($\sigma = 3$)

erence rates. As in Proposition 3, the larger the time preference rate, the smaller is the size of the overlap.

In the market economy, policy is needed to counteract the excessively short horizon of investors introduced by effect of finite patent length. A subsidy to R&D can do this job and is needed to decentralize the first-best.

OA4 Segmented labor market

Suppose labour market is segmented and the total supply of labour is $\bar{L}$ for workers and $\bar{s}$ for scientists, that is, $L_c + L_d \leq \bar{L}$ and $s_c + s_d \leq \bar{s}$, where $\bar{L}$ and $\bar{s}$ are now parameters. Wage in (8) now differs from wage in (11). Denote the former by w_L and the latter by w_s. From (11) and (24), we see that we can continue to use the variable m_j to determine the innovation regime.

Combining (12), (15), and (22), we find

$$\hat{\lambda}_j = r - (1-\alpha)\alpha \frac{\theta_j PY}{\lambda_j Q_j}. \tag{OA.63}$$

From (24), we find $\dot{m}_c = m_c(1-m_c)\left(\hat{\lambda}_c + \hat{Q}_c - \hat{\lambda}_d - \hat{Q}_d\right)$ or

$$\dot{m}_c = m_c(1-m_c)\mu(s_c - s_d) + \alpha\mu \frac{(1-\alpha)PY}{w_s} m_k(m_c - \theta_c), \tag{OA.64}$$

where k denotes the research active sector and we have used $\mu\lambda_c Q_c = w_s m_c / m_k$ based on (11). Finally, the clean market share evolves according to

$$\dot{\theta}_c = \theta_c(1-\theta_c)(\sigma - 1)\mu(s_c - s_d). \tag{OA.65}$$

From (OA.64) and (OA.65), it is clear that the two variables θ_c and m_c are insufficient in describing the entire dynamics of the model due to the expression

$(1-\alpha)PY/w_s$ in (OA.64). This expression can be rewritten as $w_L\bar{L}/w_s$. Compared to the integrated labour market, where relative wage in production and research is 1 and production labour L is endogenous, here $\bar{L}$ is a constant but the relative wage w_L/w_s is an endogenous variable. Assuming segmented labour market thus does not reduce the dimensionality of the model, because a third variable is needed in order to fully describe the dynamics of the model.

To more precisely compare the model with segmented labour market with the integrated labour market model in the main text, we define $L_e \equiv (1-\alpha)PY/w_s$. This new variable adopts a similar role to production labour L in the main text by determining the savings rate. With integrated labour market, the savings rate is given by $ws/(PC) \propto s/L = (1-L)/L$. Here, using (18), we find L_e to be inversely proportional to the savings rate:

$$L_e = \frac{\bar{s}}{1+\alpha}\left[\frac{w_s\bar{s}}{PC}\right]^{-1}.$$

Combining the above expression with (11), (OA.63) and (17), we find

$$\dot{L}_e = L_e\left[\alpha\mu L_e\theta_k - \mu\bar{s} - \rho\right]. \tag{OA.66}$$

Using L_e, (OA.64) can be written as

$$\dot{m}_c = \alpha\mu L_e m_k(m_c-\theta_c) + \mu(s_c-s_d)m_c(1-m_c). \tag{OA.67}$$

Together, (OA.65)-(OA.67) form a differential equation system in variables θ_c, L_e, and m_c that summarizes the dynamics of the model with segmented labor market. Comparing with Lemma 1, we see that the dynamics of both models are almost identical, the only difference being that where the constant $\bar{s}$ shows up in the equation for $\dot{L}_e$, in the integrated market model the variable $s_k = 1-L$ shows up.

We conclude that assuming segmented labour market does not reduce the dimensionality of the model. The reason is that, as long as the investment decision is dynamic, the expected present value of investment m_j will be affected by the savings rate. Even if the supply of scientists is fixed (measured as labor input), the savings rate (measured consumption equivalents) changes over time depending on the relative wage. If labour is mobile across the two sectors and the wage equalized, the savings rate depends on the allocation of scientists. Thus, to reduce the dimensionality, we either need to make investment decision static (e.g. with one-period patent) or assume a fixed savings rate.

OA5 Modeling innovation as creative destruction

In this appendix we model innovation as creative destruction and demonstrate that it generates the same qualitative result.

The R&D process is the following (see also Acemoglu et al., 2012). At the beginning of each period, each scientist decides whether to direct her research to clean or dirty technology. Each scientist is then randomly allocated to innovating at most one machine within their chosen sector with a success probability of μ.

The rest of the model is the same as in Section 2. Thus, equations (8)-(10), (13), (15)-(22) continue to hold. From (15), the average profit of a sector is given by

$$\pi_j = \int_0^1 \pi_{ji} di = (1-\alpha)\alpha P_j Y_j. \tag{OA.68}$$

Denote by V_j the value of the patent in sector j in case of successful innovation. The decision to target a particular sector is governed by the free entry condition:

$$\mu V_j \leq w \perp s_j \geq 0, \tag{OA.69}$$

which is the counterpart of (11). Note that for the research active sector k, $V_k = w/\mu$. From (8),

$$V_k = (1-\alpha)P_k Y_k/(\mu L_k) \tag{OA.70}$$

follows.

Equation (12) is replaced by the non arbitrage condition:

$$\dot{V}_j = rV_j - \pi_j + \mu s_j V_j. \tag{OA.71}$$

For the dynamic equilibrium, note that (31) and (32) continue to hold. Similar to (25), we can define

$$m_c \equiv \frac{V_c}{V_c + V_d}. \tag{OA.72}$$

From (OA.69), it is clear that $m_c \gtreqless 1/2$ separates the three innovation regimes. Time differentiating m_c and using (OA.71), we find

$$\dot{m}_c = m_c(1-m_c)\left(\hat{V}_c - \hat{V}_d\right) = m_c(1-m_c)\left(-\frac{\pi_c}{V_c} + \frac{\pi_d}{V_d} + \hat{Q}_c - \hat{Q}_d\right). \tag{OA.73}$$

Using (OA.68), (OA.70), (21), and the definition of m_c, we can derive (30). Thus, Lemma 1 holds.

This analysis shows that creative destruction and inhouse R&D offer almost exactly the same equilibrium conditions, both statically and dynamically. The only difference is that quality improvement is evaluated at its marginal value of improving the patent in the case of inhouse R&D, as innovation can occur repeatedly, whereas in the case of creative destruction, quality improvement is valued at the total value of the patent. As Lemma 1 holds in both cases, this difference does not matter for the result.